arXiv:1606.03463v1 [math.OC] 10 Jun 2016

Online Constrained Optimization over Time Varying Renewal Systems: An Empirical Method

Xiaohan Wei
Department of Electrical Engineering, University of Southern California
xiaohanw@usc.edu

Michael J. Neely
Department of Electrical Engineering, University of Southern California
mjneely@usc.edu

This paper considers constrained optimization over a renewal system. A controller observes a random event at the beginning of each renewal frame and then chooses an action that affects the duration of the frame, the amount of resources used, and a penalty metric. The goal is to make frame-wise decisions so as to minimize the time average penalty subject to time average resource constraints. This problem has applications to task processing and communication in data networks, as well as to certain classes of Markov decision problems. We formulate the problem as a dynamic fractional program and propose an online algorithm which adopts an empirical accumulation as a feedback parameter. Prior work considers a ratio method that needs statistical knowledge of the random events. A key feature of the proposed algorithm is that it does not require knowledge of the statistics of the random events. We prove the algorithm satisfies the desired constraints and achieves $O(\epsilon)$ near optimality with probability 1.

Key words: renewal system, Markov decision processes, stochastic optimization, online optimization
MSC2000 subject classification: 90C15, 90C30, 90C40, 93E35
History:

1. Introduction Consider a system that operates over the timeline of real numbers $t \geq 0$. The timeline is divided into back-to-back periods called *renewal frames* and the start of each frame is called a *renewal* (see Fig. 1). The system state is refreshed at each renewal. At the start of each renewal frame $n \in \{0, 1, 2, \dots\}$ the controller observes a random event $\omega[n]$ and then takes an action $\alpha[n]$ from an action set $\mathcal{A}$. The pair $(\omega[n], \alpha[n])$ affects: (i) the duration of that renewal frame; (ii) a vector of resource expenditures for that frame; (iii) a penalty incurred on that frame. The goal is to choose actions over time to minimize time average penalty subject to time average constraints on the resources.

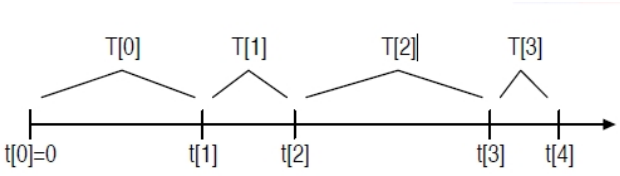

FIGURE 1. An illustration of a sequence of renewal frames.

1.1. Example applications This problem has applications to task processing and file downloading in computer networks, and to certain classes of Markov decision problems.

• Task processing: Consider a device that processes computational tasks back-to-back. Each renewal period corresponds to the time required to complete a single task. The random event $\omega[n]$ observed for task n corresponds to a vector of task parameters, including the type, size, and

resource requirements for that particular task. The action set $\mathcal{A}$ consists of different processing mode options, and the specific action $\alpha[n]$ determines the processing time, energy expenditure, and task quality. In this case, task quality can be defined as a negative penalty, and the goal is to maximize time average quality subject to power constraints and task completion rate constraints.

- File downloading: Consider a wireless device that repeatedly downloads files. The device has two states: *active* (wants to download a file) and *idle* (does not want to download a file). Here, $\omega[n]$ denotes the observed wireless channel state, which affects the success probability of downloading a file (and thereby affects the transition probability from active to idle). This example is discussed further in the simulation section (Section 6).
- Markov decision problems: Consider a discrete time Markov decision problem over an infinite horizon and with constraints on average cost per slot (see [1] and [2] for exposition of Markov decision theory). Assume there is a special state that is recurrent under any sequence of actions (similar assumptions are used in [2]). Renewals are defined as revisitation times to that state. A random event $\omega[n]$ is observed upon each revisitation and affects the Markov properties for all timeslots of that frame. This can be viewed as a *two-timescale Markov decision problem* [3][4]. In this case, the control action chosen on frame k is in fact a policy for making decisions over slots until the next renewal. The algorithm of the current paper does not require knowledge of the statistics of $\omega[n]$, but does require knowledge of the conditional Markov transition probabilities (given $\omega[n]$) over each frame n.

1.2. Prior approaches In the special case when the set of all possible random events $\omega[n]$ is finite, the action set $\mathcal{A}$ is finite, and the probabilities of $\omega[n]$ are known, the problem can be solved (offline) by finding a solution to a linear fractional program (see also [1] for basic renewal-reward theory). Methods for solving linear fractional programs are in [5][6][7]. The current paper seeks an online method that does not require statistical knowledge and that can handle possibly infinite random event sets and action sets.

Prior work in [8] considers the same renewal optimization problem as the current paper and solves it via a *drift-plus-penalty ratio* method. This method requires an action to be chosen that minimizes a ratio of expectations. However, this choice requires knowledge of the statistics of $\omega[n]$. Methods for approximate minimization of the ratio are considered in [8]. A heuristic algorithm is also proposed in [8] that is easier to implement because it does not require knowledge of statistics. That algorithm is partially analyzed: It is shown that if the process converges, then it converges to a near-optimal point. However, whether or not it converges is unknown.

1.3. Our contributions The current paper develops a new algorithm that is easy to implement, requires neither statistics of $\omega[n]$ nor explicit estimation of them, and is fully analyzed with convergence properties that provably hold with probability 1. In particular, the feasibility of the algorithm is justified through a stability analysis of virtual queues and the near optimality comes from a novel construction of exponential supermartingales. Simulation experiments on a time varying constrained MDP conform with theoretical results.

1.4. Related work The renewal system problem considered in this paper is a generalization of stochastic optimization over fixed time slots. Problems are categorized based on whether or not the random event is observed before the decisions are made. Cases where the random event is observed are useful in network optimization problems where max-weight ([9][10]), Lyapunov optimization ([11][12][13][14]), fluid model methods ([15][16]), and dual subgradient methods ([17][18][19]) are often used.

Cases where the random events are not observed are called online convex optimization. Various algorithms are developed for unconstrained learning including (but not limited to) weighted majority algorithm ([20]), multiplicative weighting algorithm ([21]), following the perturbed leader

([22]) and online gradient descent ([23][24]). The resource constrained learning problem is studied in [25] and [26]. Moreover, online learning with underlying MDP structure is also treated using modified multiplicative weighting ([27]) and improved following the perturbed leader ([28]).

2. Problem Formulation and Preliminaries

2.1. Formulation Consider a system where the time line is divided into back-to-back time periods called frames. At the beginning of frame n ($n \in \{0,1,2,\cdots\}$), a controller observes the realization of a random variable $\omega[n]$, which is an i.i.d. copy of a random variable taking values in a compact set $\Omega \in \mathbb{R}^q$ with distribution function unknown to the controller. Then, after observing the random event, the controller chooses an action vector $\alpha[n] \in \mathcal{A}$. Then, the tuple $(\omega[n],\ \alpha[n])$ induces the following random variables:

- The penalty received during frame n: $y[n]$.
- The length of frame n: $T[n]$.
- A vector of resource consumptions during frame n: $\mathbf{z}[n] = [z_1[n],\ z_2[n],\ \cdots,\ z_L[n]]$.

We assume that *given* $\alpha[n] = \alpha$ *and* $\omega[n] = \omega$ *at frame* n, $(y[n], T[n], \mathbf{z}[n])$ *is a random vector independent of the outcomes of previous frames*, with *known* expectations. We then denote these conditional expectations as

$$\begin{aligned}\hat{y}(\omega,\alpha) =& \mathbb{E}\left[y[n] \mid \omega,\alpha\right],\\ \hat{T}(\omega,\alpha) =& \mathbb{E}\left[T[n] \mid \omega,\alpha\right],\\ \hat{\mathbf{z}}(\omega,\alpha) =& \mathbb{E}\left[\hat{\mathbf{z}}[n] \mid \omega,\alpha\right],\end{aligned}$$

which are all deterministic functions of ω and α. This notation is useful when we want to highlight the action α we choose. The analysis assumes a single action in response to the observed $\omega[n]$ at each frame. Nevertheless, an ergodic MDP can fit into this model by defining the action as a selection of a policy to implement over that frame so that the corresponding $\hat{y}(\omega,\alpha)$, $\hat{T}(\omega,\alpha)$ and $\hat{\mathbf{z}}(\omega,\alpha)$ are expectations over the frame under the chosen policy.

Let

$$\begin{aligned}\overline{y}[N] &= \frac{1}{N}\sum_{n=0}^{N-1} y[n],\\ \overline{T}[N] &= \frac{1}{N}\sum_{n=0}^{N-1} T[n],\\ \overline{z}_l[N] &= \frac{1}{N}\sum_{n=0}^{N-1} z_l[n] \quad l \in \{1,2,\cdots,L\}.\end{aligned}$$

The goal is to minimize the time average penalty subject to L constraints on resource consumptions. Specifically, we aim to solve the following fractional programming problem:

$$\min \quad \limsup_{N\to\infty} \frac{\overline{y}[N]}{\overline{T}[N]} \tag{1}$$

$$\text{s.t.} \quad \limsup_{N\to\infty} \frac{\overline{z}_l[N]}{\overline{T}[N]} \le c_l, \quad \forall l \in \{1,2,\cdots,L\}, \tag{2}$$

$$\alpha[n] \in \mathcal{A},\ \forall n \in \{0,1,2,\cdots\}, \tag{3}$$

where $c_l,\ l \in \{1,2,\cdots,L\}$ are nonnegative constants, and both the minimum and constraint are taken in an almost sure sense. Finally, we use θ^* to denote the minimum that can be achieved by solving above optimization problem. For simplicity of notations, let

$$K[n] = \sqrt{\sum_{l=1}^{L} (z_l[n] - c_l T[n])^2}. \tag{4}$$

2.2. Assumptions Our main result requires the following assumptions, their importance will become clear as we proceed. We begin with the following boundedness assumption:

Assumption 1 (Exponential type) *Given $\omega[n]=\omega\in\Omega$ and $\alpha[n]=\alpha\in\mathcal{A}$ for a fixed n, it holds that $T[n]\geq 1$ with probability 1 and $y[n]$, $K[n]$, $T[n]$ are of exponential type, i.e. there exists a constant $\eta>0$ s.t.*

$$\begin{aligned}&\mathbb{E}\left[\exp\left(\eta y[n]\right)|\,\omega,\alpha\right]\leq B,\\ &\mathbb{E}\left[\exp\left(\eta K[n]\right)|\,\omega,\alpha\right]\leq B,\\ &\mathbb{E}\left[\exp\left(\eta T[n]\right)|\,\omega,\alpha\right]\leq B,\end{aligned}$$

where B is a positive constant.

The following proposition is a simple consequence of the above assumption:

Proposition 1 *Suppose Assumption 1 holds. Let $X[n]$ be any of the three random variables $y[n]$, $K[n]$ and $T[n]$ for a fixed n. Then, given any $\omega[n]=\omega\in\Omega$ and $\alpha[n]=\alpha\in\mathcal{A}$,*

$$\mathbb{E}\left[X[n]|\,\omega,\alpha\right]\leq B/\eta,\quad \mathbb{E}\left[X[n]^2\big|\,\omega,\alpha\right]\leq 2B/\eta^2.$$

The proof is easy by expanding the term $\mathbb{E}\left[e^{\eta X[n]}\big|\,\omega,\alpha\right]$ using Taylor series and bounding $\mathbb{E}\left[X[n]|\,\omega,\alpha\right]$, $\mathbb{E}\left[X[n]^2|\,\omega,\alpha\right]$ using the first and the second order term, respectively.

Assumption 2 *There exists a positive constant $\theta_{\max}$ large enough so that the optimal objective of (1) – (3), denoted as θ^*, falls into $[0,\theta_{\max})$ with probability 1.*

Remark 1 *If $\theta^*<0$, then, we shall find a constant c large enough so that $\theta^*+c\geq 0$. Then, define a new penalty $y'[n]=y[n]+cT[n]$. It is then easy to see that minimizing $\limsup_{N\to\infty}\overline{y}[N]/\overline{T}[N]$ is equivalent to minimizing $\limsup_{N\to\infty}\overline{y'}[N]/\overline{T}[N]$ and the optimal objective of the new problem is θ^*+c, which is nonnegative.*

Assumption 3 *Let $\left(\hat{y}(\omega,\alpha),\ \hat{T}(\omega,\alpha),\ \hat{\mathbf{z}}(\omega,\alpha)\right)$ be the performance vector under a certain (ω,α) pair. Then, for any fixed $\omega\in\Omega$, the set of achievable performance vectors over all $\alpha\in\mathcal{A}$ is compact.*

In order to state the next assumption, we need the notion of *randomized stationary policy*. We start with the definition:

Definition 1 (Randomized stationary policy) *A randomized stationary policy is an algorithm that at the beginning of each frame n, after observing the realization $\omega[n]$, the controller chooses $\alpha^*[n]$ with a conditional probability which depends only on $\omega[n]$.*

Assumption 4 (Bounded achievable region) *Let*

$$(\overline{y},\ \overline{T},\ \overline{\mathbf{z}})\triangleq\mathbb{E}\left[(\hat{y}(\omega[0],\alpha^*[0]),\ \hat{T}(\omega[0],\alpha^*[0]),\ \hat{\mathbf{z}}(\omega[0],\alpha^*[0]))\right]$$

be the one-shot average of one randomized stationary policy. Let $\mathcal{R}\subseteq\mathbb{R}^{L+2}$ be the set of all achievable one-shot averages $(\overline{y},\ \overline{T},\ \overline{\mathbf{z}})$. Then, $\mathcal{R}$ is bounded.

Assumption 5 (ξ-slackness) *There exists a randomized stationary policy $\alpha^{(\xi)}[n]$ such that the following holds,*

$$\frac{\mathbb{E}\left[\hat{z}_l\left(\omega[n],\alpha^{(\xi)}[n]\right)\right]}{\mathbb{E}\left[\hat{T}(\omega[n],\alpha^{(\xi)}[n])\right]}=c_l-\xi,\quad\forall l\in\{1,2,\cdots,L\},$$

where $\xi>0$ is a constant.

Remark 2 (Measurability issue) *We implicitly assume the policies for choosing α in reaction to ω result in a measurable α, so that $T[n]$, $y[n]$, $\mathbf{z}[n]$ are valid random variables and the expectations in Assumption 4 and 5 are well defined. This assumption is mild. For example, when the sets Ω and $\mathcal{A}$ are finite, it holds for any randomized stationary policy. More generally, if Ω and $\mathcal{A}$ are measurable subsets of some separable metric spaces, this holds whenever the conditional probability in Definition 1 is "regular" (see [33] for the exposition of regular conditional probability and related topics), and $T[n]$, $y[n]$, $\mathbf{z}[n]$ are continuous functions on $\Omega \times \mathcal{A}$.*

3. An Online Algorithm We define a vector of virtual queues $\mathbf{Q}[n] = [Q_1[n]\ Q_2[n]\ \cdots\ Q_L[n]]$ which are 0 at $n = 0$ and updated as follows:

$$Q_l[n+1] = \max\{Q_l[n] + z_l[n] - c_l T[n], 0\}. \tag{5}$$

The intuition behind this virtual queue idea is that if the algorithm can stabilize $Q_l[n]$, then the "arrival rate" $\overline{z}_l[N]/\overline{T}[N]$ is below "service rate" c_l and the constraint is satisfied. The proposed algorithm then proceeds as in Algorithm 1 via two fixed parameters $V > 0$, $\delta > 0$, and an additional process $\theta[n]$ that is initialized to be $\theta[0] = 0$. For any real number x, the notation $[x]_0^{\theta_{\max}}$ stands for ceil and floor function:

$$[x]_0^{\theta_{\max}} = \begin{cases} \theta_{\max}, & \text{if } x \in (\theta_{\max}, +\infty); \\ x, & \text{if } x \in [0, \theta_{\max}]; \\ 0, & \text{if } x \in (-\infty, 0). \end{cases}$$

Note that we can rewrite (6) as the following deterministic form:

$$V\left(\hat{y}(\omega[n], \alpha[n]) - \theta[n]\hat{T}(\omega[n], \alpha[n])\right) + \sum_{l=1}^{L} Q_l[n]\left(\hat{z}_l(\omega[n], \alpha[n]) - c_l\hat{T}(\omega[n], \alpha[n])\right),$$

Thus, Algorithm 1 proceeds by observing $\omega[n]$ on each frame n and then choosing $\alpha[n]$ in $\mathcal{A}$ to minimize the above deterministic function. We can now see that we only use knowledge of current realization $\omega[n]$, not statistics of $\omega[n]$. Also, the compactness assumption (Assumption 3) guarantees that the minimum of (6) is always achievable. For the rest of the paper, we introduce several abbreviations:

$$D[n] \triangleq \sum_{i=0}^{n}\left(y[i] - \theta[i]T[i] + \frac{1}{V}\sum_{l=1}^{L} Q_l[i](z_l[i] - c_l T[i])\right) \ \rightarrow \text{ empirical accumulation}$$

$$\hat{\theta}[n] \triangleq \frac{1}{(n+1)^\delta}\sum_{i=0}^{n}\left(y[i] - \theta[i]T[i] + \frac{1}{V}\sum_{l=1}^{L} Q_l[i](z_l[i] - c_l T[i])\right) \ \rightarrow \text{ original pseudo average}$$

$$\theta[n] \triangleq \left[\frac{1}{(n+1)^\delta}\sum_{i=0}^{n}\left(y[i] - \theta[i]T[i] + \frac{1}{V}\sum_{l=1}^{L} Q_l[i](z_l[i] - c_l T[i])\right)\right]_0^{\theta_{\max}} \ \rightarrow \text{ trimmed pseudo average.}$$

The intuitive reason why we need trimmed pseudo average is to ensure the empirical accumulation does not blow up, which plays an important role in the analysis.

4. Feasibility Analysis In this section, we prove that the proposed algorithm gives a sequence of actions $\{\alpha[n]\}_{n=0}^{\infty}$ which satisfies all desired constraints with probability 1. Specifically, we show that all virtual queues are stable with probability 1, in which we leverage an important lemma from [30] to obtain a exponential bound for the norm of $\mathbf{Q}[n]$.

Algorithm 1 Online renewal optimization:

- At the beginning of each frame n, the controller observes $Q_l[n]$, $\theta[n]$, $\omega[n]$ and chooses action $\alpha[n] \in \mathcal{A}$ to minimize the following function:

$$\mathbb{E}\left[V(y[n]-\theta[n]T[n]) + \sum_{l=1}^{L} Q_l[n](z_l[n]-c_l T[n])\middle| Q_l[n],\theta[n],\omega[n]\right]. \tag{6}$$

- Update $\theta[n]$:

$$\theta[n+1] = \left[\frac{1}{(n+1)^\delta}\sum_{i=0}^{n}\left(y[i]-\theta[i]T[i]+\frac{1}{V}\sum_{l=1}^{L}Q_l[i](z_l[i]-c_l T[i])\right)\right]_0^{\theta_{\max}}.$$

- Update virtual queues $Q_l[n]$:

$$Q_l[n+1] = \max\{Q_l[n]+z_l[n]-c_l T[n],0\},\ l=1,2,\cdots,L.$$

4.1. The drift-plus-penalty bound The start of our proof uses the drift-plus-penalty methodology. For a general introduction on this topic, see [29] for more details. We define the 2-norm function of the virtual queue vector as:

$$\|\mathbf{Q}[n]\|^2 = \sum_{l=1}^{L} Q_l[n]^2.$$

Define the *Lyapunov drift* $\Delta(\mathbf{Q}[n])$ as

$$\Delta(\mathbf{Q}[n]) = \frac{1}{2}\left(\|\mathbf{Q}[n+1]\|^2 - \|\mathbf{Q}[n]\|^2\right).$$

Next, define the penalty function at frame n as $V(y[n]-\theta[n]T[n])$, where $V>0$ is a fixed trade-off parameter. Then, the drift-plus-penalty methodology suggests that we can stabilize the virtual queues by choosing an action $\alpha[n] \in \mathcal{A}$ to greedily minimize the following drift-plus-penalty expression, with the observed $\mathbf{Q}[n]$, $\omega[n]$ and $\theta[n]$:

$$\mathbb{E}\left[V(y[n]-\theta[n]T[n]) + \Delta(\mathbf{Q}[n])\middle| Q_l[n],\theta[n],\omega[n]\right].$$

The penalty term $V(y[n]-\theta[n]T[n])$ uses the $\theta[n]$ variable, which depends on events from all previous frames. This penalty does not fit the rubric of [29] and convergence of the algorithm does not follow from prior work. A significant thrust of the current paper is convergence analysis via the trimmed pseudo averages defined in the previous subsection.

In order to obtain an upper bound on $\Delta(\mathbf{Q}[n])$, we square both sides of (5) and use the fact that $\max\{x,0\}^2 \le x^2$,

$$Q_l[n+1]^2 \le Q_l[n]^2 + (z_l[n]-c_l T[n])^2 + 2Q_l[n](z_l[n]-c_l T[n]). \tag{7}$$

Then we have

$$\begin{aligned}
&\mathbb{E}\left[V(y[n]-\theta[n]T[n]) + \Delta(\mathbf{Q}[n])\middle| Q_l[n],\theta[n],\omega[n]\right]\\
\le&\mathbb{E}\left[V(y[n]-\theta[n]T[n]) + \sum_{l=1}^{L} Q_l[n](z_l[n]-c_l T[n])\middle| Q_l[n],\theta[n],\omega[n]\right] + \frac{1}{2}\sum_{l=1}^{L}\mathbb{E}\left[(z_l[n]-c_l T[n])^2\right]\\
\le&\mathbb{E}\left[V(y[n]-\theta[n]T[n]) + \sum_{l=1}^{L} Q_l[n](z_l[n]-c_l T[n])\middle| Q_l[n],\theta[n],\omega[n]\right] + \frac{B^2}{\eta^2}.
\end{aligned} \tag{8}$$

where the last inequality follows from Proposition 1. Thus, as we have already seen in Algorithm 1, the proposed algorithm observes the vector $\mathbf{Q}[n]$, the random event $\omega[n]$ and the trimmed pseudo average $\theta[n]$ at frame n, and minimizes the right hand side of (8).

4.2. Bounds on the virtual queue process and feasibility In this section, we show how the bound (8) leads to the feasibility of the proposed algorithm. Define $\mathcal{H}_n$ as the system history information up until frame n. Formally, $\{\mathcal{H}_n\}_{n=0}^{\infty}$ is a filtration where each $\mathcal{H}_n$ is the σ-algebra generated by all the random variables before frame n. Notice that since $\mathbf{Q}[n]$ and $\theta[n]$ depend only on the events before frame n, $\mathcal{H}_n$ contains both $\mathbf{Q}[n]$ and $\theta[n]$. The following important lemma gives a stability criterion for any given real random process with certain negative drift property:

Lemma 1 (Theorem 2.3 of [30]) *Let $R[n]$ be a real random process over $n \in \{0,1,2,\cdots\}$ satisfying the following two conditions for a fixed $r>0$:*

1. *For any n, $\mathbb{E}\left[e^{r(R[n+1]-R[n])}\middle|\mathcal{H}_n\right] \le \Gamma$, for some $\Gamma>0$.*
2. *Given $R[n] \ge \sigma$, $\mathbb{E}\left[e^{r(R[n+1]-R[n])}\middle|\mathcal{H}_n\right] \le \rho$, with some $\rho \in (0,1)$.*

Suppose further that $R[0] \in \mathbb{R}$ is given and finite, then, at every $n \in \{0,1,2,\cdots\}$, the following bound holds:

$$\mathbb{E}\left[e^{rR[n]}\right] \le \rho^n e^{rR[0]} + \frac{1-\rho^n}{1-\rho}\Gamma e^{r\sigma}.$$

Thus, in order to show the stability of the virtual queue process, it is enough to test the above two conditions with $R[n] = \|\mathbf{Q}[n]\|$. The following lemma shows that $\|\mathbf{Q}[n]\|$ satisfies these two conditions:

Lemma 2 (Drift condition) *Let $R[n] = \|\mathbf{Q}[n]\|$, then, it satisfies the two conditions in Lemma 1 with the following constants:*

$$\begin{aligned}
&\Gamma = B,\\
&r = \min\left\{\eta, \frac{\xi\eta^2}{4B}\right\},\\
&\sigma = C_0 V,\\
&\rho = 1 - \frac{r\xi}{2} + \frac{2B}{\eta^2}r^2 < 1.
\end{aligned}$$

where $C_0 = \frac{2B^2}{V\xi\eta^2} + \frac{2(\theta_{\max}+1)B}{\xi\eta} - \frac{\xi}{4V}$.

The central idea of the proof is to plug the ξ-slackness policy specified in Assumption 5 into the right hand side of (8). A similar idea has been presented in the Lemma 6 of [31] under the bounded increment of the virtual queue process. Here, we generalize the idea to the case where the increment of the virtual queues contains exponential type random variables $z_l[n]$ and $T[n]$. Note that the boundedness of $\theta[n]$ is crucial for the argument to hold, which justifies the truncation of pseudo average in the algorithm. Lemma 1 is proved in the Appendix.

Combining the above two lemmas, we immediately have the following corollary:

Corollary 1 (Exponential decay) *Given $\mathbf{Q}[0] = 0$, the following holds for any $n \in \{0,1,2,\cdots\}$ under the proposed algorithm,*

$$\mathbb{E}\left[e^{r\|\mathbf{Q}[n]\|}\right] \le D, \tag{9}$$

where

$$D = 1 + \frac{B}{1-\rho}e^{rC_0V},$$

and r, ρ, C_0 are as defined in Lemma 2.

With Corollary 1 in hand, we can prove the following theorem:

Theorem 1 (Feasibility) *All constraints in (1)-(3) are satisfied under the proposed algorithm with probability 1.*

Proof of Theorem 1. By queue updating rule (5), for any n and any $l \in \{1, 2, \cdots, L\}$, one has

$$Q_l[n+1] \geq Q_l[n] + z_l[n] - c_l T[n].$$

Fix N as a positive integer. Then, summing over all $n \in \{0, 1, 2, \cdots, N-1\}$,

$$Q_l[N] \geq Q_l[0] + \sum_{n=0}^{N-1} (z_l[n] - c_l T[n]).$$

Since $Q_l[0] = 0, \ \forall l$ and $T[n] \geq 1, \ \forall n$,

$$\frac{\sum_{n=0}^{N-1} z_l[n]}{\sum_{n=0}^{N-1} T[n]} - c_l \leq \frac{Q_l[N]}{\sum_{n=0}^{N-1} T[n]} \leq \frac{Q_l[N]}{N}. \tag{10}$$

Define the event

$$A_N^{(\varepsilon)} = \{Q_l[N] > \varepsilon N\}.$$

By the Markov inequality and Corollary 1, for any $\varepsilon > 0$, we have

$$\begin{aligned} Pr(Q_l[N] > \varepsilon N) \leq & Pr\left(r\|\mathbf{Q}[N]\| > r\varepsilon N\right) \\ = & Pr\left(e^{r\|\mathbf{Q}[N]\|} > e^{r\varepsilon N}\right) \\ \leq & \frac{\mathbb{E}\left[e^{r\|\mathbf{Q}[N]\|}\right]}{e^{r\varepsilon N}} \leq De^{-r\varepsilon N}, \end{aligned}$$

where r is defined in Corollary 1. Thus, we have

$$\sum_{N=0}^{\infty} Pr(Q_l[N] > \varepsilon N) \leq D \sum_{N=0}^{\infty} e^{-r\varepsilon N} < +\infty.$$

Thus, by the Borel-Cantelli lemma ([33]),

$$Pr\left(A_N^{(\varepsilon)} \text{ occurs infinitely often}\right) = 0.$$

Since $\varepsilon > 0$ is arbitrary, letting $\varepsilon \to 0$ gives

$$Pr\left(\lim_{N\to\infty} \frac{Q_l[N]}{N} = 0\right) = 1.$$

Finally, taking the $\limsup_{N\to\infty}$ from both sides of (10) and substituting in the above equation gives the claim. □

5. Optimality Analysis In this section, we show that the proposed algorithm achieves time average penalty within $\mathcal{O}(1/V)$ of the optimal objective θ^*. Since the algorithm meets all the constraints, it follows,

$$\limsup_{n\to\infty} \frac{\sum_{i=0}^{n-1} y[i]}{\sum_{i=0}^{n-1} T[i]} \geq \theta^*, \quad w.p.1.$$

Thus, it is enough to prove the following theorem:

Theorem 2 (Near optimality) *For any $\delta \in (1/3, 1)$ and $V \geq 1$, the objective value produced by the proposed algorithm is near optimal with*

$$\limsup_{n\to\infty} \frac{\sum_{i=0}^{n-1} y[i]}{\sum_{i=0}^{n-1} T[i]} \leq \theta^* + \frac{B^2}{\eta^2 V}, \ w.p.1,$$

i.e. the algorithm achieves $\mathcal{O}(1/V)$ near optimality.

The proof of this theorem is relatively involved, thus, we would like to sketch the roadmap of our proof before jumping into the details.

The key point of the proof is to bound the pseudo average $\hat{\theta}[n]$ asymptotically from above by θ^*, which is achieved in Theorem 3 below. We then prove Theorem 3 through the following three-step-constructions:

- Introduce a term-wise truncated version of $\hat{\theta}[n]$, denoted as $\tilde{\theta}[n]$, who has the same limit as $\hat{\theta}[n]$ (shown in Lemma 5), so that it is enough to show $\tilde{\theta}[n] \leq \theta^*$ asymptotically.
- Construct a process $F[n]$ as in (14) and show that $F[n] \geq 0$ is a necessary condition for $\tilde{\theta}[n] > \theta^*$ (shown in Lemma 7).
- Show that $\tilde{\theta}[n] > \theta^*$ only finitely often (with probability 1) by an exponential supermartingale construction related to $F[n]$ (shown in Lemma 8 and Lemma 9).

5.1. Relation between $\hat{\theta}[n]$ and $\theta[n]$ We start with a preparation lemma illustrating that the original pseudo average $\hat{\theta}[n]$ behaves almost the same as the trimmed pseudo average $\theta[n]$. Recall that $\theta[n]$ is defined as:

$$\theta[n] = [\hat{\theta}[n]]_0^{\theta_{max}}.$$

Lemma 3 (Equivalence relation) *For any $x \in (0, \theta_{\max})$,*

1. *$\theta[n] \geq x$ if and only if $\hat{\theta}[n] \geq x$.*
2. *$\theta[n] \leq x$ if and only if $\hat{\theta}[n] \leq x$.*
3. *$\limsup_{n\to\infty} \theta[n] \leq x$ if and only if $\limsup_{n\to\infty} \hat{\theta}[n] \leq x$.*
4. *$\limsup_{n\to\infty} \theta[n] \geq x$ if and only if $\limsup_{n\to\infty} \hat{\theta}[n] \geq x$.*

This lemma is intuitive and the proof is shown in the Appendix. We will see that this $\hat{\theta}[n]$ is sometimes easier to work with than $\theta[n]$, and we will prove results on $\hat{\theta}[n]$ which extend naturally to $\theta[n]$ by Lemma 3.

5.2. Towards near optimality (I): Truncation The following lemma states that the optimality of (1)-(3) is achievable within the closure of the set of all one-shot averages specified in Assumption 4:

Lemma 4 (Stationary optimality) *Let θ^* be the optimal objective of (1)-(3). Then, there exists a tuple $(y^*, T^*, \mathbf{z}^*) \in \overline{\mathcal{R}}$, the closure of $\mathcal{R}$, such that the following hold:*

$$y^*/T^* = \theta^* \tag{11}$$
$$z_l^*/T^* \leq c_l, \ \forall l \in \{1, 2, \cdots, L\}, \tag{12}$$

i.e. the optimality is achievable within $\overline{\mathcal{R}}$.

The proof of this lemma is similar to the proof of Theorem 4.5 as well as Lemma 7.1 of [32]. We omit the details for brevity.

We start the truncation by picking up an $\varepsilon_0 > 0$ small enough so that $\theta^* + \varepsilon_0/V < \theta_{\max}$. We aim to show $\limsup_{n\to\infty} \theta[n] \leq \theta^* + \varepsilon_0/V$. By Lemma 3, it is enough to show $\limsup_{n\to\infty} \hat{\theta}[n] \leq \theta^* + \varepsilon_0/V$. The following lemma tells us it is enough to prove it on a further term-wise truncated version of $\hat{\theta}[n]$.

Lemma 5 (Truncation lemma) *Consider the following alternative pseudo average* $\{\tilde{\theta}[n]\}_{n=0}^{\infty}$ *by truncating each summand such that* $\tilde{\theta}[0]=0$ *and*

$$\tilde{\theta}[n+1]=\frac{1}{(n+1)^{\delta}}\sum_{i=0}^{n}\left[\left(y[i]-\theta[i]T[i]+\frac{1}{V}\sum_{l=1}^{L}Q_l[i](z_l[i]-c_lT[i])\right)\wedge\left(\left(\frac{2}{\eta}+\frac{4\sqrt{L}}{\eta rV}\right)\log^2(i+1)\right)\right],$$

where $a\wedge b\triangleq\min\{a,b\}$, η *is defined in Assumption 1 and* r *is defined in Lemma 2. Then, we have*

$$\limsup_{n\to\infty}\hat{\theta}[n]=\limsup_{n\to\infty}\tilde{\theta}[n].$$

Proof of Lemma 5. Consider any frame i such that there is a discrepancy between the summand of $\hat{\theta}[n]$ and $\tilde{\theta}[n]$, i.e.

$$y[i]-\theta[i]T[i]+\frac{1}{V}\sum_{l=1}^{L}Q_l[i](z_l[i]-c_lT[i])>\left(\frac{2}{\eta}+\frac{4\sqrt{L}}{\eta rV}\right)\log^2(i+1),\tag{13}$$

By Cauchy-Schwartz inequality, this implies

$$y[i]-\theta[i]T[i]+\frac{1}{V}\sqrt{\sum_{l=1}^{L}Q_l[i]^2}\sqrt{\sum_{l=1}^{L}(z_l[i]-c_lT[i])^2}>\left(\frac{2}{\eta}+\frac{4\sqrt{L}}{\eta rV}\right)\log^2(i+1).$$

Thus, at least one of the following three events happened:

1. $A_i\triangleq\left\{y[i]-\theta[i]T[i]>\frac{2}{\eta}\log^2(i+1)\right\}$.
2. $B_i\triangleq\left\{\sqrt{\sum_{l=1}^{L}Q_l[i]^2}>\frac{2\sqrt{L}}{r}\log(i+1)\right\}$.
3. $E_i\triangleq\left\{K[i]>\frac{2}{\eta}\log(i+1)\right\}$.

where $K[i]$ is defined in (4). Indeed, the occurence of one of the three events is necessary for (13) to happen. We then argue that these three events jointly occur only finitely many times. Thus, as $n\to\infty$, the discrepancies are negligible.

Assume the event A_i occurs, then, since $y[i]-\theta[i]T[i]\le y[i]$, it follows $y[i]>\frac{2}{\eta}\log^2(i+1)$. Then, we have

$$\begin{aligned}Pr(A_i)\le&Pr\left(y[i]>\frac{2}{\eta}\log^2(i+1)\right)\\=&Pr\left(e^{\eta y[i]}>e^{2\log^2(i+1)}\right)\\\le&\frac{\mathbb{E}\left[e^{\eta y[i]}\right]}{(i+1)^{2\log(i+1)}}\le\frac{B}{(i+1)^{2\log(i+1)}},\end{aligned}$$

where the second to last inequality follows from Markov inequality and the last inequality follows from Assumption 1.

Assume the event B_i occurs, then, we have

$$\|\mathbf{Q}[i]\|=\sqrt{\sum_{l=1}^{L}Q_l[i]^2}>\frac{2\sqrt{L}}{r}\log(i+1)\ge\frac{2}{r}\log(i+1).$$

Thus,

$$\begin{aligned}Pr(B_i)\le&Pr\left(\|\mathbf{Q}[i]\|>\frac{2}{r}\log(i+1)\right)\\=&Pr\left(e^{r\|\mathbf{Q}[i]\|}>e^{2\log(i+1)}\right)\\\le&\frac{\mathbb{E}\left[e^{r\|\mathbf{Q}[i]\|}\right]}{(i+1)^2}\le\frac{D}{(i+1)^2},\end{aligned}$$

where the second to last inequality follows from Markov inequality and the last inequality follows from Corollary 1.

Assume the event E_i occurs. Again, by Assumption 1 and Markov inequality,

$$\begin{aligned} Pr(E_i) =& Pr\left(K[i] > \frac{2}{\eta}\log(i+1)\right) \\ =& Pr\left(e^{\eta K[i]} > e^{2\log(i+1)}\right) \\ \leq& \frac{\mathbb{E}\left[e^{\eta K[i]}\right]}{(i+1)^2} \leq \frac{B}{(i+1)^2}, \end{aligned}$$

where the last inequality follows from Assumption 1 again. Now, by a union bound,

$$Pr(A_i \cup B_i \cup E_i) \leq Pr(A_i) + Pr(B_i) + Pr(E_i) \leq \frac{B}{(i+1)^{2\log(i+1)}} + \frac{B+D}{(i+1)^2},$$

and thus,

$$\sum_{i=0}^{\infty} Pr(A_i \cup B_i \cup E_i) \leq \sum_{i=0}^{\infty}\left(\frac{B}{(i+1)^{2\log(i+1)}} + \frac{B+D}{(i+1)^2}\right) < \infty$$

By the Borel-Cantelli lemma, we have the joint event $A_i \cup B_i \cup E_i$ occurs only finitely many times with probability 1, and our proof is finished. □

Lemma 5 is crucial for the rest of the proof. Specifically, it creates an alternative sequence $\tilde{\theta}[n]$ which has the following two properties:

1. We know exactly what the upper bound of each of the summands is, whereas in $\hat{\theta}[n]$, there is no exact bound for the summand due to $Q_l[i]$ and other exponential type random variables.
2. For any $n \in \mathbb{N}$, we have $\tilde{\theta}[n] \leq \hat{\theta}[n]$. Thus, if $\tilde{\theta}[n] \geq \theta^* + \varepsilon_0/V$ for some n, then, $\hat{\theta}[n] \geq \theta^* + \varepsilon_0/V$.

5.3. Towards near optimality (II): Exponential supermartingale The following preliminary lemma demonstrates a negative drift property for each of the summands in $\tilde{\theta}[n]$.

Lemma 6 (Key feature inequality) *For any $\varepsilon_0 > 0$, if $\theta[i] \geq \theta^* + \varepsilon_0/V$, then, we have*

$$\mathbb{E}\left[\left(y[i] - \theta[i]T[i] + \frac{1}{V}\sum_{l=1}^{L} Q_l[i](z_l[i] - c_l T[i])\right) \wedge \left(\left(\frac{2}{\eta} + \frac{4\sqrt{L}}{\eta r V}\right)\log^2(i+1)\right) \middle| \mathcal{H}_i\right] \leq -\varepsilon_0/V,$$

Proof of Lemma 6. Since the proposed algorithm minimizes (6) over all possible decisions in $\mathcal{A}$, it must achieve value less than or equal to that of any randomized stationary algorithm $\alpha^*[i]$. This in turn implies,

$$\begin{aligned} &\mathbb{E}\left[\left(y[i] - \theta[i]T[i] + \frac{1}{V}\sum_{l=1}^{L} Q_l[i](z_l[i] - c_l T[i])\right) \middle| \mathcal{H}_i, \omega[i]\right] \\ \leq &\mathbb{E}\left[\left(\hat{y}(\omega[i], \alpha^*[i]) - \theta[i]\hat{T}(\omega[i], \alpha^*[i]) + \frac{1}{V}\sum_{l=1}^{L} Q_l[i](\hat{z}_l(\omega[i], \alpha^*[i]) - c_l\hat{T}(\omega[i], \alpha^*[i]))\right) \middle| \mathcal{H}_i, \omega[i]\right]. \end{aligned}$$

Taking expectation from both sides with respect to $\omega[i]$ and using the fact that randomized stationary algorithms are i.i.d. over frames and independent of $\mathcal{H}_i$, we have

$$\mathbb{E}\left[\left(y[i] - \theta[i]T[i] + \frac{1}{V}\sum_{l=1}^{L} Q_l[i](z_l[i] - c_l T[i])\right) \middle| \mathcal{H}_i, \omega[i]\right] \leq \overline{y} - \theta[i]\overline{T} + \frac{1}{V}\sum_{l=1}^{L} Q_l[i](\overline{z}_l - c_l\overline{T}),$$

for any $(\overline{y}, \overline{T}, \overline{\mathbf{z}}) \in \mathcal{R}$. Since $(y^*, T^*, \mathbf{z}^*)$ specified in Lemma 4 is in the closure of $\mathcal{R}$, we can replace $(\overline{y}, \overline{T}, \overline{\mathbf{z}})$ by the tuple $(y^*, T^*, \mathbf{z}^*)$ and the inequality still holds. This gives

$$\begin{aligned}
&\mathbb{E}\left[\left(y[i]-\theta[i]T[i]+\frac{1}{V}\sum_{l=1}^{L}Q_l[i](z_l[i]-c_lT[i])\right)\middle|\mathcal{H}_i,\omega[i]\right]\\
\leq& y^*-\theta[i]T^*+\frac{1}{V}\sum_{l=1}^{L}Q_l[i](z_l^*-c_lT^*),\\
=&T^*\left(y^*/T^*-\theta[i]+\frac{1}{V}\sum_{l=1}^{L}Q_l[i](z_l^*/T^*-c_l)\right)\\
\leq& T^*(\theta^*-\theta[i])\leq -\varepsilon_0/V,
\end{aligned}$$

where the second to last inequality follows from (11) and (12), and the last inequality follows from $\theta[i] \geq \theta^* + \varepsilon_0/V$ and $T[i] \geq 1$. Finally, since $a \wedge b \leq a$ for any real numbers a, b, it follows,

$$\begin{aligned}
&\mathbb{E}\left[\left(y[i]-\theta[i]T[i]+\frac{1}{V}\sum_{l=1}^{L}Q_l[i](z_l[i]-c_lT[i])\right)\wedge\left(\left(\frac{2}{\eta}+\frac{4\sqrt{L}}{\eta r V}\right)\log^2(i+1)\right)\middle|\mathcal{H}_i\right]\\
\leq&\mathbb{E}\left[\left(y[i]-\theta[i]T[i]+\frac{1}{V}\sum_{l=1}^{L}Q_l[i](z_l[i]-c_lT[i])\right)\middle|\mathcal{H}_i\right]\leq -\varepsilon_0/V,
\end{aligned}$$

and the claim follows. □

Define n_k as the frame where $\tilde{\theta}[n]$ visits the set $(-\infty,\ \theta^* + \varepsilon_0/V)$ for the k-th time with the following conventions: 1. If $\tilde{\theta}[n] \in (-\infty,\ \theta^* + \varepsilon_0/V)$ and $\tilde{\theta}[n+1] \in (-\infty,\ \theta^* + \varepsilon_0/V)$, then we count them as two times. 2. When $k = 1$, n_1 is equal to 0. Define the *hitting time* S_{n_k} as

$$S_{n_k} = n_{k+1} - n_k.$$

The goal is to obtain a moment bound on this quantity when $\tilde{\theta}[n_k + 1] \geq \theta^* + \varepsilon_0/V$ (otherwise, this quantity is 1). In order to do so, we introduce a new process as follows. For any n_k, define

$$F[n] \triangleq \sum_{i=n_k}^{n-1}\left(y[i]-\theta[i]T[i]+\frac{1}{V}\sum_{l=1}^{L}Q_l[i](z_l[i]-c_lT[i])\right)\wedge\left(\left(\frac{2}{\eta}+\frac{4\sqrt{L}}{\eta r V}\right)\log^2(i+1)\right),\ \forall n > n_k, \tag{14}$$

The following lemma shows that indeed this $F[n]$ is closely related to $\tilde{\theta}[n]$. It plays an important role in proving both Lemma 8 and Lemma 9:

Lemma 7 *For any $n > n_k$, if $\tilde{\theta}[n] \geq \theta^* + \varepsilon_0/V$, then, $F[n] \geq 0$.*

Proof of Lemma 7. Suppose $\tilde{\theta}[n] \geq \theta^* + \varepsilon_0/V$, then, the following holds

$$\theta^* + \varepsilon_0/V \leq \tilde{\theta}[n] = \frac{n_k^\delta}{n^\delta}\tilde{\theta}[n_k] + \frac{1}{n^\delta}F[n].$$

Thus,

$$F[n] \geq n^\delta(\theta^* + \varepsilon_0/V) - n_k^\delta\tilde{\theta}[n_k].$$

Since at the frame n_k, $\tilde{\theta}[n_k] < \theta^* + \varepsilon_0/V$, it follows,

$$F[n] \geq \left(n^\delta - n_k^\delta\right)(\theta^* + \varepsilon_0/V).$$

Since $\theta^* + \varepsilon_0/V \geq 0$, it follows $F[n] \geq 0$ and the claim follows. □

Recall our goal is to bound the hitting time S_{n_k} of the process $\tilde{\theta}[n]$ when $\{\tilde{\theta}[n_k+1] \geq \theta^* + \varepsilon_0/V\}$, with a strictly negative drift property as Lemma 6. A classical approach analyzing the hitting time of a stochastic process came from Wald's construction of martingale for sequential analysis (see [34] for details). Later, Hajek extended this idea to analyze the stability of a queueing system with drift condition by a supermartingale construction ([30]). Here, we take one step further by considering the following supermartingale construction related to $F[n]$:

Lemma 8 (Exponential Supermartingale) *Fix $\varepsilon_0 > 0$ and $V \geq \max\left\{\frac{\varepsilon_0\eta}{4\log^2 2} - \frac{2\sqrt{L}}{r},\ 1\right\}$ such that $\theta^* + \varepsilon_0/V < \theta_{\max}$. Define a new random process $G[n]$ starting from n_k+1 with*

$$G[n] \triangleq \frac{\exp\left(\lambda_n F[n \wedge (n_k + S_{n_k})]\right)}{\prod_{i=n_k+1}^{n\wedge(n_k+S_{n_k})} \rho_i} \mathbf{1}_{\{\tilde{\theta}[n_k+1] \geq \theta^* + \varepsilon_0/V\}},$$

where for any set A, $\mathbf{1}_A$ is the indicator function which takes value 1 if A is true and 0 otherwise. For any $n \geq n_k+1$, λ_n and ρ_n are defined as follows:

$$\lambda_n = \frac{\varepsilon_0}{2Ve\left(\frac{2}{\eta} + \frac{4\sqrt{L}}{\eta r V}\right)^2 \log^4(n+1)},$$
$$\rho_n = 1 - \frac{\varepsilon_0^2}{4V^2 e\left(\frac{2}{\eta} + \frac{4\sqrt{L}}{\eta r V}\right)^2 \log^4(n+1)}.$$

Then, the process $G[n]$ is measurable with respect to $\mathcal{H}_n$, $\forall n \geq n_k+1$, and furthermore, it is a supermartingale with respect to the filtration $\{\mathcal{H}_n\}_{n \geq n_k+1}$.

The proof of Lemma 8 is shown in the Appendix.

Remark 3 *If the increments $F[n+1] - F[n]$ were to be bounded, then, we could adopt the similar construction as that of [30]. However, in our scenario $F[n+1] - F[n]$ is of the order $\log^2(n+1)$, which is increasing and unbounded. Thus, we need decreasing exponents λ_n and increasing weights ρ_n to account for that. Furthermore, the indicator function indicates that we are only interested in the scenario $\{\tilde{\theta}[n_k+1] \geq \theta^* + \varepsilon_0/V\}$.*

The following lemma uses the previous result to bound the fourth moment of the hitting time S_{n_k} the proof is also given in the Appendix. The proof is heavy in computation mainly because the time varying exponents and weights we introduced in the last Lemma.

Lemma 9 *Given $V \geq \max\left\{\frac{\varepsilon_0\eta}{4\log^2 2} - \frac{2\sqrt{L}}{r},\ 1\right\}$ as in Lemma 8, for any $\beta \in (0, 1/5)$ and any $\varepsilon_0 > 0$ such that $\theta^* + \varepsilon_0/V < \theta_{\max}$, there exists a positive constant $C_{\beta,V,\varepsilon_0} \simeq \mathcal{O}\left(V^{10}\beta^{-20}\varepsilon_0^{-10}\right)$, such that*

$$\mathbb{E}\left[S_{n_k}^4 | \mathcal{H}_{n_k}\right] \leq C_{\beta,V,\varepsilon_0}(n_k+2)^{4\beta}, \quad \forall k \geq 1.$$

5.4. Asymptotic upper bound for $\theta[n]$ So far, we have proved that if we pick any $\varepsilon_0 > 0$ such that $\theta^* + \varepsilon_0/V < \theta_{\max}$, then, the inter-visiting time has bounded conditional fourth moment. We aim to show that $\limsup_{n\to\infty} \hat{\theta}[n] \leq \theta^*$ with probability 1. By Lemma 5, it is enough to show $\limsup_{n\to\infty} \tilde{\theta}[n] \leq \theta^*$. To do so, we need the following Second Borel-Cantelli lemma:

Lemma 10 (Theorem 5.3.2. of [33]) *Let $\mathcal{F}_k$, $k \geq 1$ be a filtration with $\mathcal{F}_1 = \{\emptyset, \Omega\}$, and A_k, $k \geq 1$ be a sequence of events with $A_k \in \mathcal{F}_{k+1}$, then*

$$\{A_k \text{ occurs infinitely often}\} = \left\{\sum_{k=1}^{\infty} Pr(A_k|\mathcal{F}_k) = \infty\right\}$$

Theorem 3 (Asymptotic upper bound) *For any $\delta \in (1/3, 1)$ and $V \geq 1$, the following hold,*

$$\limsup_{n\to\infty} \hat{\theta}[n] \leq \theta^*, \quad w.p.1,$$

and

$$\limsup_{n\to\infty} \theta[n] \leq \theta^*, \quad w.p.1.$$

Proof of Theorem 3. First of all, since the inter-hitting time S_{n_k} has finite fourth moment, each inter-hitting time is finite with probability 1, and thus the process $\{\tilde{\theta}[n]\}_{n=0}^{\infty}$ will visit $(-\infty, \theta^* + \varepsilon_0/V)$ infinitely many times with probability 1. Then, we pick any $\epsilon > 0$ and define the following sequence of events:

$$A_k \triangleq \left\{ \frac{S_{n_k}}{n_k^{1/3}} > \epsilon \right\}, \ k = 1, 2, \cdots. \tag{15}$$

For any fixed k, by Conditional Markov inequality, the following holds with probability 1:

$$\begin{aligned} Pr(A_k|\mathcal{H}_{n_k}) =& Pr\left(S_{n_k}^4 > \epsilon^4 n_k^{4/3} \middle| \mathcal{H}_{n_k} \right) \\ \leq& \frac{\mathbb{E}\left[S_{n_k}^4 | \mathcal{H}_{n_k} \right]}{\epsilon^4 n_k^{4/3}} \\ \leq& \frac{C_{\beta,V,\varepsilon_0}(n_k+2)^{4\beta}}{\epsilon^4 n_k^{4/3}} \\ \leq& \frac{C_{\beta,V,\varepsilon_0}}{\epsilon^4} n_k^{-4/3+4\beta} + \frac{C_{\beta,V,\varepsilon_0} 2^{4\beta}}{\epsilon^4 n_k^{4/3}} \\ \leq& \frac{C_{\beta,V,\varepsilon_0}}{\epsilon^4} k^{-4/3+4\beta} + \frac{C_{\beta,V,\varepsilon_0} 2^{4\beta}}{\epsilon^4} k^{-4/3}, \end{aligned}$$

where the second inequality follows from Lemma 9 with $\beta \in (0, 1/5)$, the third inequality follows from the fact that $(a+b)^x \leq a^x + b^x, \ \forall a, b \geq 0$ and $x \in (0,1)$. The last inequality follows from the fact that the inter-hitting time takes at least one frame and thus $n_k \geq k$.

Choose $\mathcal{F}_k = \mathcal{H}_{n_k}$ and A_k as is defined in (15). Then, for any $\beta \in (0, 1/12)$, we have with probability 1,

$$\sum_{k=1}^{\infty} Pr(A_k|\mathcal{H}_{n_k}) \leq \sum_{k=1}^{\infty} \left(\frac{C_{\beta,V,\varepsilon_0}}{\epsilon^4} k^{-4/3+4\beta} + \frac{C_{\beta,V,\varepsilon_0} 2^{4\beta}}{\epsilon^4} k^{-4/3} \right) < \infty.$$

Now by Lemma 10,

$$Pr\left(A_k \text{ occurs infinitely often} \right) = 0.$$

Since the process $\{\tilde{\theta}[n]\}_{n=0}^{\infty}$ visits $(-\infty, \theta^* + \varepsilon_0/V)$ infinitely many times with probability 1,

$$\limsup_{n\to\infty} \frac{S_{n_k}}{n_k^{1/3}} = \limsup_{k\to\infty} \frac{S_{n_k}}{n_k^{1/3}} \leq \epsilon, \ w.p.1,$$

Since $\epsilon > 0$ is arbitrary, let $\epsilon \to 0$ gives

$$\lim_{n\to\infty} \frac{S_{n_k}}{n_k^{1/3}} = 0, \ w.p.1. \tag{16}$$

Finally, we show how this convergence result leads to the bound of $\tilde{\theta}[n]$. According to the updating rule of $\tilde{\theta}[n]$, for any frame n such that $n_k < n \le n_{k+1}$,

$$\begin{aligned}\tilde{\theta}[n] =&(\frac{n_k}{n})^\delta \tilde{\theta}[n_k] + \frac{1}{n^\delta}\sum_{i=n_k}^{n-1}\left(y[i]-\theta[i]T[i]+\frac{1}{V}Q[i](z[i]-cT[i])\right)\wedge\left(\left(\frac{2}{\eta}+\frac{4\sqrt{L}}{\eta r V}\right)\log^2(i+1)\right)\\ \le&(\frac{n_k}{n})^\delta\left(\theta^*+\frac{\varepsilon_0}{V}\right)+\frac{1}{n^\delta}\sum_{i=n_k}^{n-1}\left(\left(\frac{2}{\eta}+\frac{4\sqrt{L}}{\eta r V}\right)\log^2(i+1)\right)\\ \le&(\frac{n_k}{n})^\delta\left(\theta^*+\frac{\varepsilon_0}{V}\right)+\frac{1}{n^\delta}S_{n_k}\left(\frac{2}{\eta}+\frac{4\sqrt{L}}{\eta r V}\right)\log^2 n,\end{aligned}$$

where the first inequality follows from the fact that $\tilde{\theta}[n_k] < \theta^* + \varepsilon_0/V$. Now, we take the $\limsup_{n\to\infty}$ from both sides and analyze each single term on the right hand side:

$$\begin{aligned}&1 \ge \limsup_{n\to\infty}(\frac{n_k}{n})^\delta \ge \limsup_{k\to\infty}(\frac{n_k}{n_k+S_{n_k}})^\delta = \limsup_{k\to\infty}(\frac{1}{1+\frac{S_{n_k}}{n_k}})^\delta = 1,\ w.p.1,\\ &\limsup_{n\to\infty}\frac{S_{n_k}}{n^\delta}\left(\frac{2}{\eta}+\frac{4\sqrt{L}}{\eta r V}\right)\log^2 n \le \limsup_{n\to\infty}\frac{S_{n_k}}{n_k^{1/3}}\cdot\limsup_{n\to\infty}\frac{\left(\frac{2}{\eta}+\frac{4\sqrt{L}}{\eta r V}\right)\log^2 n}{n^{\delta-1/3}} = 0,\ w.p.1,\end{aligned}$$

where we apply the convergence result (16) in the second line. Thus,

$$\limsup_{n\to\infty}\tilde{\theta}[n] \le \theta^* + \frac{\varepsilon_0}{V},\ w.p.1.$$

By Lemma 5 we have $\limsup_{n\to\infty}\hat{\theta}[n] \le \theta^* + \varepsilon_0/V$. Finally, by Lemma 3, and the fact that $\theta^* + \varepsilon_0/V \in (0, \theta_{\max})$, we have $\limsup_{n\to\infty}\theta[n] \le \theta^* + \varepsilon_0/V$. Since this holds for any $\varepsilon_0 > 0$ small enough, let $\varepsilon_0 \to 0$ finishes the proof. □

5.5. Finishing the proof of near optimality With the help of previous analysis on $\theta[n]$, we are ready to prove our main theorem, with the following lemma on strong law of large numbers for martingale difference sequences:

Lemma 11 (Corollary 4.2 of [29]) *Let $\{\mathcal{F}_i\}_{i=0}^\infty$ be a filtration and let $\{X(i)\}_{i=0}^\infty$ be a real-valued random process such that $X(i) \in \mathcal{F}_{i+1},\ \forall i$. Suppose there is a finite constant C such that $\mathbb{E}[X(i)|\mathcal{F}_i] \le C,\ \forall i$, and*

$$\sum_{i=1}^\infty \frac{\mathbb{E}[X(i)^2]}{i^2} < \infty.$$

Then,

$$\limsup_{n\to\infty}\frac{1}{n}\sum_{i=0}^{n-1}X(i) \le C,\quad w.p.1.$$

Proof of Theorem 2. Recall for any n, the empirical accumulation without ceil and floor function is

$$\hat{\theta}[n] = \frac{1}{n^\delta}\sum_{i=0}^{n-1}\left(y[i]-\theta[i]T[i]+\frac{1}{V}\sum_{l=1}^{L}Q_l[i](z_l[i]-c_lT[i])\right).$$

Dividing both sides by $\sum_{i=0}^{n-1} T[i]/n^\delta$ yields

$$\begin{aligned}\frac{\hat{\theta}[n]}{\frac{1}{n^\delta}\sum_{i=0}^{n-1} T[i]} &= \frac{\sum_{i=0}^{n-1}\left(y[i]-\theta[i]T[i]+\frac{1}{V}\sum_{l=1}^{L} Q_l[i](z_l[i]-c_l T[i])\right)}{\sum_{i=0}^{n-1} T[i]} \\ &= \frac{\sum_{i=0}^{n-1}\left(y[i]+\frac{1}{V}\sum_{l=1}^{L} Q_l[i](z_l[i]-c_l T[i])\right)}{\sum_{i=0}^{n-1} T[i]} - \frac{\sum_{i=0}^{n-1}\theta[i]T[i]}{\sum_{i=0}^{n-1} T[i]}.\end{aligned}$$

Moving the last term to the left hand side and taking the $\limsup_{n\to\infty}$ from both sides gives

$$\begin{aligned}\limsup_{n\to\infty}\left(\frac{\hat{\theta}[n]}{\frac{1}{n^\delta}\sum_{i=0}^{n-1} T[i]}+\frac{\sum_{i=0}^{n-1}\theta[i]T[i]}{\sum_{i=0}^{n-1} T[i]}\right) &\geq \limsup_{n\to\infty}\frac{\sum_{i=0}^{n-1} y[i]}{\sum_{i=0}^{n-1} T[i]}+\frac{\sum_{i=0}^{n-1}\frac{1}{V}\sum_{l=1}^{L} Q_l[i](z_l[i]-c_l T[i])}{\sum_{i=0}^{n-1} T[i]} \\ &\geq \limsup_{n\to\infty}\frac{\sum_{i=0}^{n-1} y[i]}{\sum_{i=0}^{n-1} T[i]}+\frac{1}{2}\frac{\|\mathbf{Q}[n]\|^2-\sum_{i=0}^{n-1}\sum_{l=1}^{L}(z_l[i]-c_l T[i])^2}{V\sum_{i=0}^{n-1} T[i]} \\ &\geq \limsup_{n\to\infty}\frac{\sum_{i=0}^{n-1} y[i]}{\sum_{i=0}^{n-1} T[i]}-\frac{1}{2V}\limsup_{n\to\infty}\frac{1}{n}\sum_{i=0}^{n-1} K[i]^2,\end{aligned}$$

where the second inequality follows from inequality (7) and telescoping sums, and the last inequality follows from $T[n]\geq 1$, $\|\mathbf{Q}[n]\|^2\geq 0$ and $K[i]=\sqrt{\sum_{l=1}^{L}(z_l[i]-c_l T[i])^2}$. Now we use Lemma 11 with $X(i)=K[i]^2$ to bound the second term. Since $K[i]$ is of exponential type by Assumption 1, we know that $\mathbb{E}[K[i]^2|\mathcal{H}_n]\leq 2B^2/\eta^2$. Furthermore, $\mathbb{E}[K[i]^4]\leq 24B^4/\eta^4$. Thus,

$$\sum_{i=1}^{\infty}\frac{\mathbb{E}[K[i]^4]}{i^2}<\infty.$$

Thus, all assumptions in Lemma 11 are satisfied and we conclude that

$$\limsup_{n\to\infty}\frac{1}{n}\sum_{i=0}^{n-1} K[i]^2\leq\frac{2B^2}{\eta^2},\ w.p.1.$$

This implies,

$$\limsup_{n\to\infty}\left(\frac{\hat{\theta}[n]}{\frac{1}{n^\delta}\sum_{i=0}^{n-1} T[i]}+\frac{\sum_{i=0}^{n-1}\theta[i]T[i]}{\sum_{i=0}^{n-1} T[i]}\right)\geq\limsup_{n\to\infty}\frac{\sum_{i=0}^{n-1} y[i]}{\sum_{i=0}^{n-1} T[i]}-\frac{B^2}{\eta^2 V}.$$

By Theorem 3, $\hat{\theta}[n]$ is asymptotically upper bounded. Since $\delta<1$ and $T[n]\geq 1$, it follows $\frac{1}{n^\delta}\sum_{i=0}^{n-1} T[i]=\mathcal{O}(n^{1-\delta})$, which goes to infinity as $n\to\infty$. Thus,

$$\limsup_{n\to\infty}\frac{\hat{\theta}[n]}{\frac{1}{n^\delta}\sum_{i=0}^{n-1} T[i]}\leq 0,$$

and thus,

$$\limsup_{n\to\infty}\frac{\sum_{i=0}^{n-1}\theta[i]T[i]}{\sum_{i=0}^{n-1} T[i]}\geq\limsup_{n\to\infty}\frac{\sum_{i=0}^{n-1} y[i]}{\sum_{i=0}^{n-1} T[i]}-\frac{B^2}{\eta^2 V}.$$

By Theorem 3 again, $\theta[n]$ is asymptotically upper bounded by θ^*. Based on this result, it is easy to show the following

$$\limsup_{n\to\infty}\frac{\sum_{i=0}^{n-1}\theta[i]T[i]}{\sum_{i=0}^{n-1} T[i]}\leq\theta^*.$$

Thus, we finally get

$$\limsup_{n\to\infty}\frac{\sum_{i=0}^{n-1} y[i]}{\sum_{i=0}^{n-1} T[i]}\leq\theta^*+\frac{B^2}{\eta^2 V},$$

finishing the proof. □

6. Simulation experiments In this section, we demonstrate the performance of our proposed algorithm through an application scenario on single user file downloading. We show that this problem can be formulated as a two state constrained online MDP and solved using our proposed algorithm.

Consider a slotted time system where $t \in \{0,1,2,\cdots\}$, and one user is repeatedly downloading files. We use $F(t) \in \{0,1\}$ to denote the system file state at time slot t. State "1" indicates there is an active file in the system for downloading and state "0" means there is no file and the system is idle. Suppose the user can only download 1 file at each time, and the user cannot observe the file length. Each file contains an integer number of packets which is independent and geometrically distributed with expected length equal to 1.

During each time slot where there is an active file for downloading (i.e. $F(t)=1$), the user first observes the channel state $\omega(t)$, which is the i.i.d. random variable taking values in $\Omega = \{0.2, 0.5, 0.8\}$ with equal probabilities, and delay penalty $s(t)$, which is also an i.i.d. random variable taking values in $\{1,3,5\}$ with equal probability. Then, the user makes a service action $\alpha(t) \in \mathcal{A} = \{0, 0.3, 0.6, 0.9\}$. The pair $(\omega(t), \alpha(t))$ affects the following quantities:

- The success probability of downloading a file at time t: $\phi(\alpha(t), \omega(t)) \triangleq \alpha(t) \cdot \omega(t)$.
- The resource consumption $p(\alpha(t))$ at time t. We assume $p(0)=0$, $p(0.3)=1$, $p(0.6)=2$ and $p(0.9)=4$.

After a file is downloaded, the system goes idle (i.e. $F(t)=0$) and stays there for a random amount of time that is independent and geometrically distributed with mean equal to 2. The goal is to minimize the time average delay penalty subject to a resource constraint that the time average resource consumption cannot exceed 1.

In [35], a similar optimization problem is considered but without random events $\omega(t)$ and $s(t)$, which can be formulated as a two state constrained MDP. Here, using the same logic, we can formulate our optimization problem as a two state constrained online MDP. Given $F(t)=1$, the file will finish its download at the end of this time slot with probability $\phi(\alpha(t), \omega(t))$. Thus, the transition probabilities out of state 1 are:

$$Pr[F(t+1)=0|F(t)=1] = \phi(\alpha(t), \omega(t))$$
$$Pr[F(t+1)=1|F(t)=1] = 1 - \phi(\alpha(t), \omega(t)),$$

On the other hand, given $F(t)=0$, the system is idle and will transition to the active state in the next slot with probability λ:

$$Pr[F(t+1)=1|F(t)=0] = \lambda$$
$$Pr[F(t+1)=0|F(t)=0] = 1 - \lambda,$$

Now, we characterize this online MDP through renewal frames and show that it can be solved using the proposed algorithm in Section 2. First, notice that the state "1" is recurrent under any action $\alpha(t)$. We denote t_n as the n-th time slot when the system returns to state "1". Define the renewal frame as the time period between t_n and t_{n+1} with frame size

$$T[n] = t_{n+1} - t_n.$$

Furthermore, since the system does not have any control options in state "0", the controller makes exactly one decision during each frame and this decision is made at the beginning of each frame. Thus, we can write out the optimization problem as follows:

$$\begin{aligned} \min \quad & \limsup_{N\to\infty} \frac{\sum_{n=0}^{N-1} \alpha(t_n) s(t_n)}{\sum_{n=0}^{N-1} T[n]} \\ s.t. \quad & \limsup_{N\to\infty} \frac{\sum_{n=0}^{N-1} p(\alpha(t_n))}{\sum_{n=0}^{N-1} T[n]} \leq 1, \ \alpha(t_n) \in \mathcal{A}. \end{aligned}$$

Subsequently, in order to apply our algorithm, we can define the virtual queue $Q[n]$ as $Q[0]=0$ with updating rule

$$Q[n+1]=\max\{Q[n]+p(\alpha(t_n))-T[n],0\}.$$

Notice that for any particular action $\alpha(t_n)\in\mathcal{A}$ and random event $\omega(t_n)\in\Omega$, we can always compute $\mathbb{E}[T[n]]$ as

$$\begin{aligned}\mathbb{E}[T[n]]&=1-\phi(\alpha(t_n),\omega(t_n))+\phi(\alpha(t_n),\omega(t_n))\left(1+\frac{1}{\lambda}\right)\\&=1+2\alpha(t_n)\omega(t_n),\end{aligned}$$

where the second equality follows by substituting $\lambda=0.5$ and $\phi(\alpha(t_n),\omega(t_n))=\alpha(t_n)\omega(t_n)$. Thus, for each $\alpha(t_n)\in\mathcal{A}$, the expression (6) can be computed.

In each of the simulations, each data point is the time average of 2 million slots. We compare the performance of the proposed algorithm with the optimal randomized policy. The optimal policy is computed by formulating the MDP into a linear program with the knowledge of the distribution on $\omega(t)$ and $s(t)$. See [6] for details of this linear program formulation.

In Fig. 2, we plot the performance of our algorithm verses V parameter for different δ value. We see from the plots that as V gets larger, the time averages approaches the optimal value and achieves a near optimal performance for δ roughly between 0.4 and 1. A more obvious relation between performance and δ value is shown in Fig. 3, where we fix $V=300$ and plot the performance of the algorithm verses δ value. It is clear from the plots that the algorithm fails whenever δ is too small ($\delta<0.3$) or too big ($\delta>1$). This meets the statement of Theorem 2 that the algorithm works for $\delta\in(1/3,1)$.

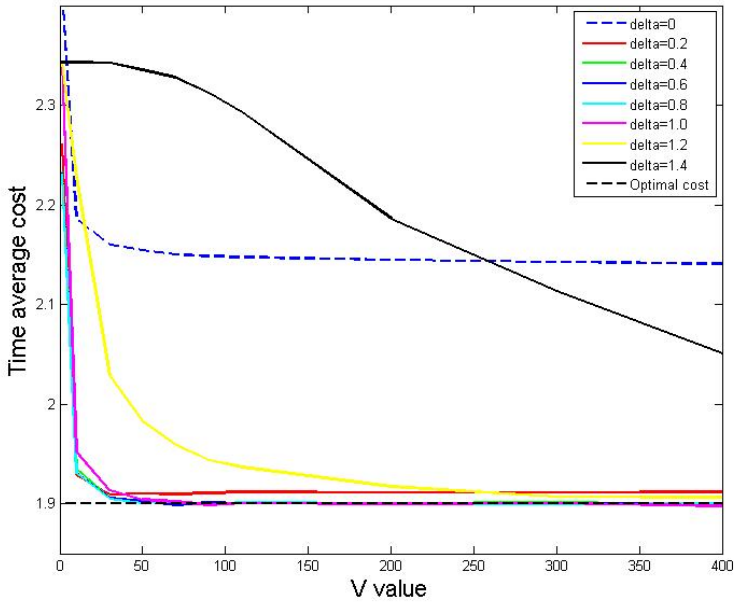

FIGURE 2. Time average penalty versus tradeoff parameter V

In Fig. 4, we plot the time average resource consumption verses V value. We see from the plots that the algorithm is always feasible for different V's and δ's, which meets the statement of Theorem 1. Also, as V gets larger, the constraint gap tends to be smaller. In Fig. 5, we plot the average

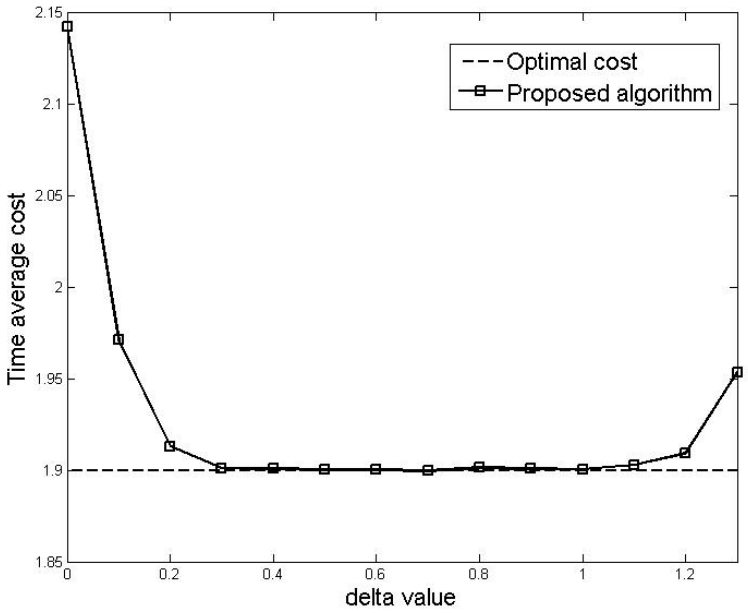

FIGURE 3. Time average penalty versus δ parameter with fixed $V = 300$.

virtual queue size verses V value. It shows that the average queue size gets larger as V get larger. To see the implications, recall from the proof of Theorem 1, the inequality (10) implies that the virtual queue size $Q_l[N]$ affects the rate that the algorithm converges down to the feasible region. Thus, if the average virtual queue size is large, then, it takes longer for the algorithm to converge. This demonstrates that V is indeed a trade-off parameter which trades the sub-optimality gap for the convergence rate.

7. Conclusions This paper considers the constrained optimization over a renewal system with observed random events at the beginning of each renewal frame. We propose an online algorithm which does not need the knowledge of the distributions of random events. We prove that this proposed algorithm is feasible and achieves $O(\varepsilon)$ near optimality by constructing an exponential supermartingale. Simulation experiments demonstrates the near optimal performance of the proposed algorithm.

References

[1] S. Ross. 2002. *Introduction to Probability Models, 8th edition.* Academic Press.

[2] D. P. Bertsekas. 2001. *Dynamic Programming and Optimal Control, 2nd edition, Vol. I.* Athena Scientific, Nashua, NH.

[3] C. Wernz. 2013. Multi-time-scale Markov decision processes for organizational decision making. *EURO J. Decis. Process,* **1**, pp. 299-324.

[4] H. S. Chang, P. J. Fard, S. I. Marcus, and M. Shayman. 2003. Multi-time Scale Markov Decision Processes. *IEEE Trans. Automatic Control,* **48**(6), pp. 976-987.

[5] S. Schaible. 1983. Fractional programming. *Zeitschrift fur Operations Research,* **27**(1), pp. 39-54.

[6] B. Fox. 1966. Markov renewal programming by linear fractional programming. *SIAM Journal on Applied Mathematics*, **14**(6), pp. 1418-1432.

[7] S. Boyd and L. Vandenberghe. 2004. *Convex Optimization.* Cambridge University Press.

[8] M. J. Neely. 2013. Dynamic optimization and learning for renewal systems. *IEEE Transactions on Automatic Control,* **58**(1), pp. 32-46.

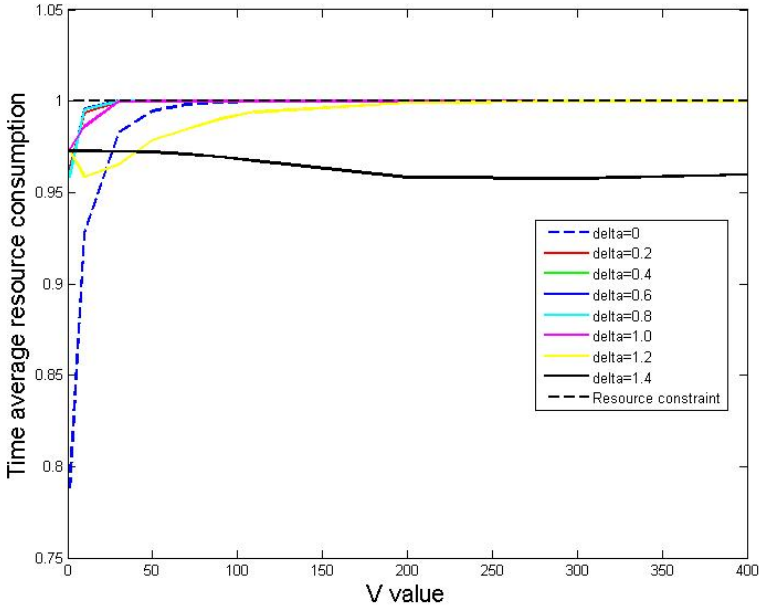

FIGURE 4. Time average resource consumption versus tradeoff parameter V.

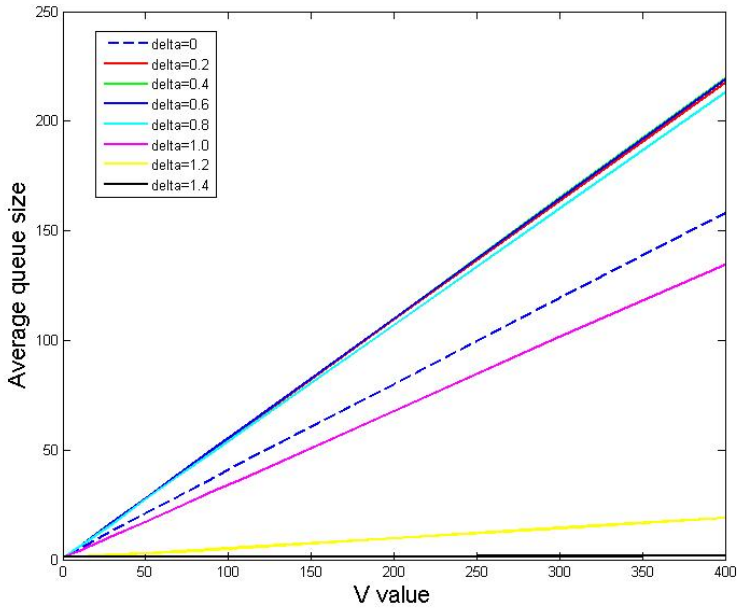

FIGURE 5. Time average virtual queue size versus tradeoff parameter V.

[9] L. Tassiulas and A. Ephremides. 1992. Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks. *IEEE Transacations on Automatic Control,* **37**(12), pp. 1936-1948.

[10] L. Tassiulas and A. Ephremides. 1993. Dynamic server allocation to parallel queues with randomly varying connectivity. *IEEE Transactions on Information Theory,* **39**(2), pp. 466-478.

[11] A. Eryilmaz and R. Srikant. 2006. *Joint congestion control, routing, and MAC for stability and fairness in wireless networks. IEEE Journal on Selected Areas in Communications, Special Issue on Nonlinear Optimization of Communication Systems,* **14**, pp. 1514-1524.

[12] A. Eryilmaz and R. Srikant. 2007. Fair resource allocation in wireless networks using queue-length-based scheduling and congestion control. *IEEE/ACM Transactions on Networking,* **15**(6), pp. 1333-1344.

[13] M. J. Neely. 2010. *Stochastic Network Optimization with Application to Communication and Queueing Systems.* Morgan& Claypool.

[14] L. Georgiadis, M. J. Neely, and L. Tassiulas. 2006. Resource allocation and cross-layer control in wireless networks. *Foundations and Trends in Networking,* **1**(1), pp. 1-149.

[15] A. Stolyar. 2005. Maximizing queueing network utility subject to stability: Greedy primal-dual algorithm. *Queueing Systems,* **50**(4), pp. 401-457.

[16] A. Eryilmaz, R. Srikant. 2007. Fair resource allocation in wireless networks using queue-length-based scheduling and congestion control. *IEEE/ACM Transactions on Networking*, **15**(6), pp. 1333-1344.

[17] X. Lin and N. B. Shroff. 2006. The impact of imperfect scheduling on cross-layer congestion control in wireless networks. *IEEE/ACM Transactions on Networking,* **14**(2), pp. 302-315.

[18] X. Lin and N. B. Shroff. 2004. Joint rate control and scheduling in multihop wireless networks. *Proc. of 43rd IEEE Conf. on Decision and Control,* Paradise Island, Bahamas.

[19] A. Ribeiro. 2011. Ergodic Stochastic Optimization Algorithms for Wireless Communication and Networking. *IEEE Transactions on Signal Processing*, **58**(12), pp. 6369-6386.

[20] N. Littlestone, M. K. Warmuth. 1994. The weighted majority algorithm. *Inform. Comput.*, **108**(2), pp. 221-261.

[21] Y. Freund, R. E. Schapire. 1999. Adaptive game playing using multiplicative weights. *Games Econom. Behav.*, **29**(12), pp. 79-103.

[22] M. Hutter, J. Poland. 2005. Adaptive online prediction by following the perturbed leader. *Journal of Machine Learning Research*, **6**, pp. 639-660.

[23] M. Zinkevich. 2003. Online convex programming and generalized infinitesimal gradient ascent. *Proc. Twentieth Internat. Conf. Machine Learning.*, AAAI Press, Cambridge, MA.

[24] Elad Hazan and Satyen Kale. 2014. Beyond the regret minimization barrier: optimal algorithms for stochastic strongly-convex optimization. *The Journal of Machine Learning Research,* **15**(1), pp. 2489-2512.

[25] M. Mahdavi, R. Jin, T. Yang. 2012. Trading regret for efficiency: Online convex optimization with long term constraints. *Journal of Machine Learning Research*, **13**, pp. 2465-2490.

[26] H. Wu, R. Srikant, X. Liu, C. Jiang. 2015. Algorithm with logarithmic or sublinear regret for constrained contextual bandits. *Advances in Neural Information Processing Systems(NIPS) 28.*, Montreal, CA.

[27] Even-Dar, E., S. Kakade, Y. Mansour. 2004. Experts in a Markov decision process. *Advances in Neural Information Processing Systems(NIPS) 17.*, MIT Press, Cambridge, 401-408.

[28] J. Y. Yu, S. Mannor, N. Shimkin. 2009. Markov decision processes with arbitrary reward processes. Mathematics of Operations Research, **34**, pp. 737-757.

[29] M. J. Neely. 2012. Stability and probability 1 convergence for queueing networks via Lyapunov optimization. *Journal of Applied Mathematics*, doi:10.1155/2012/831909.

[30] B. Hajek. 1982. Hitting-time and occupation-time bounds implied by drift analysis with applications. *Advances in Applied Probability*, **14**(3), pp. 502-525.

[31] X. Wei, H. Yu, M. J. Neely. 2015. A sample path convergence time analysis of drift-plus-penalty for stochastic optimization *arXiv technical report, arXiv:1510.02973.*

[32] M. J. Neely. 2010. *Stochastic Network Optimization with Application to Communication and Queueing Systems.* Morgan & Claypool.

[33] R. Durrett. 2013. *Probability: Theory and Examples, 4th edition*, Cambridge University Press.

[34] A. Wald. 1944. On cumulative sums of random variables. *Annals of Mathematical Statistics*, **15**, pp. 283-296.

[35] X. Wei, H. Yu, M. J. Neely. 2015. Power-aware wireless file downloading: A Lyapunov indexing approach to a constrained restless bandit problem. *IEEE/ACM Transactions on Networking*, doi: 10.1109/TNET.2015.2456933.

Appendix A: Proofs

Proof of Lemma 2. We begin by bounding the difference $|\|\mathbf{Q}[n+1]\| - \|\mathbf{Q}[n]\||$ for any n:

$$\begin{aligned}
&|\|\mathbf{Q}[n+1]\| - \|\mathbf{Q}[n]\||\\
\leq&\left|\sqrt{\sum_{l=1}^{L}(Q_l[n]+z_l[n]-c_lT[n])^2}-\sqrt{\sum_{l=1}^{L}Q_l[n]^2}\right|\\
\leq&\|\mathbf{Q}[n]\|+\sqrt{\sum_{l=1}^{L}(z_l[n]-c_lT[n])^2}-\|\mathbf{Q}[n]\|\\
=&\sqrt{\sum_{l=1}^{L}(z_l[n]-c_lT[n])^2}=K[n].
\end{aligned}$$

Thus, it follows,

$$|\mathbb{E}\left[\,\|\mathbf{Q}[n+1]\|-\|\mathbf{Q}[n]\|\,|\,\mathcal{H}_n\right]|\leq\mathbb{E}\left[\,K[n]|\,\mathcal{H}_n\right]\leq\frac{B}{\eta},$$

which follows from Proposition 1. Also, we have

$$\begin{aligned}
\mathbb{E}\left[e^{r(\|\mathbf{Q}[n+1]\|-\|\mathbf{Q}[n]\|)}\middle|\,\mathcal{H}_n\right]\leq&\mathbb{E}\left[\exp\left(rK[n]\right)|\,\mathcal{H}_n\right]\\
\leq&\mathbb{E}\left[\exp\left(\eta K[n]\right)|\,\mathcal{H}_n\right]\leq B\triangleq\Gamma
\end{aligned}$$

where the second to last inequality follows by substituting the definition $r=\min\left\{\eta,\frac{\xi\eta^2}{4B}\right\}\leq\eta$ and the last inequality follows from Assumption 1.

Next, suppose $\|\mathbf{Q}[n]\|>\sigma\triangleq C_0V$. Then, since the proposed algorithm minimizes the term on the right hand side of (8) over all possible decisions at frame n, it must achieve smaller value on that term compared to that of ξ-slackness policy $\alpha^{(\xi)}[n]$ specified in Assumption 5. Formally, this is

$$\begin{aligned}
&\mathbb{E}\left[\sum_{l=1}^{L}Q_l[n](z_l[n]-c_lT[n])+V(y[n]-\theta[n]T[n])\;\middle|\;\mathcal{H}_n,\omega[n]\right]\\
\leq&\mathbb{E}\left[\sum_{l=1}^{L}Q_l[n](z_l^{(\xi)}[n]-c_lT^{(\xi)}[n])+V(y^{(\xi)}[n]-\theta[n]T^{(\xi)}[n])\;\middle|\;\mathcal{H}_n,\omega[n]\right].
\end{aligned}$$

where we used the fact that $\theta[n]$ and $\mathbf{Q}[n]$ are in $\mathcal{H}_n$. Substitute this bound into the right hand side of (8) and take expectation from both sides regarding $\omega[n]$ gives

$$\begin{aligned}
&\mathbb{E}\left[\Delta[n]+V(y[n]-\theta[n]T[n])\;|\;\mathcal{H}_n\right]\\
\leq&\mathbb{E}\left[\sum_{l=1}^{L}Q_l[n](z_l^{(\xi)}[n]-c_lT^{(\xi)}[n])+V(y^{(\xi)}[n]-\theta[n]T^{(\xi)}[n])\;\middle|\;\mathcal{H}_n\right]+B^2/\eta^2.
\end{aligned}$$

Since $\Delta[n]=\frac{1}{2}(\|\mathbf{Q}[n+1]\|^2-\|\mathbf{Q}[n]\|^2)$, This implies

$$\begin{aligned}
&\mathbb{E}\left[\|\mathbf{Q}[n+1]\|^2-\|\mathbf{Q}[n]\|^2\;|\;\mathcal{H}_n\right]\\
\leq&2B^2/\eta^2+2\mathbb{E}\left[\sum_{l=1}^{L}Q_l[n](z_l^{(\xi)}[n]-c_lT^{(\xi)}[n])+V(y^{(\xi)}[n]-\theta[n]T^{(\xi)}[n])-V(y[n]-\theta[n]T[n])\middle|\,\mathcal{H}_n\right]\\
\leq&2B^2/\eta^2+2\sum_{l=1}^{L}Q_l[n]\mathbb{E}\left[z_l^{(\xi)}[n]-c_lT^{(\xi)}[n]\middle|\,\mathcal{H}_n\right]+2V\frac{B+\theta_{\max}B}{\eta}
\end{aligned}$$

$$\leq 2B^2/\eta^2 + 2V\frac{B+\theta_{\max}B}{\eta} - 2\xi\sum_{l=1}^{L} Q_l[n]$$
$$\leq 2B^2/\eta^2 + 2V\frac{B+\theta_{\max}B}{\eta} - 2\xi\|\mathbf{Q}[n]\|,$$

where the second inequality follows from applying Proposition 1 to bound $\mathbb{E}\left[T[n]|\mathcal{H}_n\right]$ as well as the fact that $0<\theta[n]<\theta_{\max}$, and the third inequality follows from the ξ-slackness property as well as the assumption that $z_l^{(\xi)}[n]$ is i.i.d. over slots and hence independent of $Q_l[n]$. This further implies

$$\begin{aligned}
&\mathbb{E}\left[\|\mathbf{Q}[n+1]\|^2 \mid \mathcal{H}_n\right]\\
\leq&\|\mathbf{Q}[n]\|^2 - 2\xi\|\mathbf{Q}[n]\| + 2B^2/\eta^2 + 2V\frac{B+\theta_{\max}B}{\eta}\\
=&\|\mathbf{Q}[n]\|^2 - 2\xi\|\mathbf{Q}[n]\| + 2B^2/\eta^2 + 2V\frac{B+\theta_{\max}B}{\eta} - \frac{\xi^2}{4} + \frac{\xi^2}{4}\\
=&\|\mathbf{Q}[n]\|^2 - 2\xi\|\mathbf{Q}[n]\| + \frac{2B^2/\eta^2 + 2V\frac{B+\theta_{\max}B}{\eta} - \frac{\xi^2}{4}}{\xi}\cdot\xi + \frac{\xi^2}{4}\\
=&\|\mathbf{Q}[n]\|^2 - 2\xi\|\mathbf{Q}[n]\| + C_0V\cdot\xi + \frac{\xi^2}{4}\\
\leq&\|\mathbf{Q}[n]\|^2 - \xi\|\mathbf{Q}[n]\| + \frac{\xi^2}{4} = \left(\|\mathbf{Q}[n]\| - \frac{\xi}{2}\right)^2,
\end{aligned}$$

where we use the fact that $C_0 = \frac{2B^2}{V\xi\eta^2} + \frac{2}{\xi}\frac{B+\theta_{\max}B}{\eta} - \frac{\xi}{4V}$ and also the assumption that $\|\mathbf{Q}[n]\| \geq C_0V$. Now take the square root from both sides gives

$$\sqrt{\mathbb{E}\left[\|\mathbf{Q}[n+1]\|^2 \mid \mathcal{H}_n\right]} \leq \|\mathbf{Q}[n]\| - \frac{\xi}{2}.$$

By concavity of $\sqrt{x}$ function, we have $\mathbb{E}\left[\|\mathbf{Q}[n+1]\| \mid \mathcal{H}_n\right] \leq \sqrt{\mathbb{E}\left[\|\mathbf{Q}[n+1]\|^2 \mid \mathcal{H}_n\right]}$, thus,

$$\mathbb{E}\left[\|\mathbf{Q}[n+1]\| \mid \mathcal{H}_n\right] \leq \|\mathbf{Q}[n]\| - \frac{\xi}{2}. \tag{17}$$

Finally, we claim that this gives that under the condition $\|\mathbf{Q}[n]\| > \sigma \triangleq C_0V$,

$$\mathbb{E}\left[e^{r(\|\mathbf{Q}[n+1]\| - \|\mathbf{Q}[n]\|)}\middle|\mathcal{H}_n\right] \leq \rho \triangleq 1 - \frac{r\xi}{2} + \frac{2B}{\eta^2}r^2 < 1. \tag{18}$$

To see this, we expand $\mathbb{E}\left[e^{r(\|\mathbf{Q}[n+1]\| - \|\mathbf{Q}[n]\|)}\middle|\mathcal{H}_n\right]$ using Taylor series as follows:

$$\begin{aligned}
&\mathbb{E}\left[e^{r(\|\mathbf{Q}[n+1]\| - \|\mathbf{Q}[n]\|)}\middle|\mathcal{H}_n\right]\\
=&1 + r\mathbb{E}\left[\|\mathbf{Q}[n+1]\| - \|\mathbf{Q}[n]\|\middle|\mathcal{H}_n\right] + r^2\sum_{k=2}^{\infty}\frac{r^{k-2}\mathbb{E}\left[(\|\mathbf{Q}[n+1]\| - \|\mathbf{Q}[n]\|)^k\middle|\mathcal{H}_n\right]}{k!}\\
\leq&1 - \frac{r\xi}{2} + r^2\sum_{k=2}^{\infty}\frac{r^{k-2}\mathbb{E}\left[(\|\mathbf{Q}[n+1]\| - \|\mathbf{Q}[n]\|)^k\middle|\mathcal{H}_n\right]}{k!}\\
\leq&1 - \frac{r\xi}{2} + r^2\sum_{k=2}^{\infty}\frac{\eta^{k-2}\mathbb{E}\left[(\|\mathbf{Q}[n+1]\| - \|\mathbf{Q}[n]\|)^k\middle|\mathcal{H}_n\right]}{k!}\\
=&1 - \frac{r\xi}{2} + r^2\frac{\left(\mathbb{E}\left[e^{\eta(\|\mathbf{Q}[n+1]\| - \|\mathbf{Q}[n]\|)}\middle|\mathcal{H}_n\right] - \eta\mathbb{E}\left[\|\mathbf{Q}[n+1]\| - \|\mathbf{Q}[n]\|\middle|\mathcal{H}_n\right] - 1\right)}{\eta^2}\\
\leq&1 - \frac{r\xi}{2} + \frac{B+\eta\cdot\frac{B}{\eta}}{\eta^2}r^2\\
\leq&1 - \frac{r\xi}{2} + \frac{2B}{\eta^2}r^2 = \rho,
\end{aligned}$$

where the first inequality follows from (17), the second inequality follows from $r \le \eta$, and the second to last inequality follows from Proposition 1.

Finally, notice that the above quadratic function on r attains the minimum at the point $r = \frac{\xi\eta^2}{4B}$ with value $1 - \frac{\xi^2\eta^2}{8B} < 1$, and this function is strictly decreasing when

$$r \in \left(0, \frac{\xi\eta^2}{4B}\right).$$

Thus, our choice of

$$r = \min\left\{\eta, \frac{\xi\eta^2}{4B}\right\} \le \frac{\xi\eta^2}{4B}$$

ensures that ρ is strictly less than 1 and the proof is finished. □

Proof of Lemma 3. If $\theta[n] = y$ for some $y \in [0, \theta_{\max}]$, then, $\hat{\theta}[n]$ falls into one of the following three cases:

- $\hat{\theta}[n] = y$.
- $y = \theta_{\max}$ and $\hat{\theta}[n] > \theta_{\max}$.
- $y = 0$ and $\hat{\theta}[n] < 0$.

Then, we prove the above four properties based on these three cases.

1) If $\theta[n] = y \ge x$ for some y, then, the first two cases immediately imply $\hat{\theta}[n] \ge x$. If $y = 0$, then, we have $x \le 0$, which violates the assumption that $x \in (0, \theta_{\max})$. Thus, the third case is ruled out. On the other hand, if $\hat{\theta}[n] \ge x$, then, obviously, $\theta[n] \ge x$.

2) If $\theta[n] = y \le x$ for some y, then the last two cases immediately imply $\hat{\theta}[n] \le x$. If $y = \theta_{\max}$, then, we have $x \ge y_{\max}$, which violates the assumption that $x \in (0, \theta_{\max})$. Thus, the first case is ruled out. On the other hand, if $\hat{\theta}[n] \le x$, then, obviously, $\theta[n] \le x$.

3) If $\limsup_{n\to\infty} \theta[n] \le x$, then, for any $\epsilon > 0$ such that $x + \epsilon < y_{\max}$, there exists an N large enough so that $\theta[n] \le x + \epsilon$, $\forall n \ge N$. Then, by property 2), $\hat{\theta}[n] \le x + \epsilon$, $\forall n \ge N$, which implies $\limsup_{n\to\infty} \hat{\theta}[n] \le x + \epsilon$. Let $\epsilon \to 0$ gives $\limsup_{n\to\infty} \hat{\theta}[n] \le x$. One the other hand, if $\limsup_{n\to\infty} \hat{\theta}[n] \le x$, then, obviously, $\limsup_{n\to\infty} \theta[n] \le x$.

4) If $\liminf_{n\to\infty} \theta[n] \ge x$, then, for any $\epsilon > 0$ such that $x - \epsilon > 0$ there exists an N large enough so that $\theta[n] \ge x - \epsilon$, $\forall n \ge N$. Then, by property 1), $\hat{\theta}[n] \ge x - \epsilon$, $\forall n \ge N$, which implies $\limsup_{n\to\infty} \hat{\theta}[n] \le x - \epsilon$. Let $\epsilon \to 0$ gives $\limsup_{n\to\infty} \hat{\theta}[n] \ge x$. One the other hand, if $\limsup_{n\to\infty} \hat{\theta}[n] \le x$, then, obviously, $\limsup_{n\to\infty} \theta[n] \le x$. □

Proof of Lemma 8. The proof is divided into two parts. The first part contains some technical subtleties, and the second part contains computations to prove the supermartingale claim.

- *Technical preliminaries:* First of all, for any fixed k, since n_k is a random variable on the integers, it is not immediately clear that $\{\mathcal{H}_n\}_{n \ge n_k+1}$ is indeed a filtration. In order to justify this, we need the precise meaning of $\mathcal{H}_n$ for any $n \ge n_k + 1$. First, it is obvious that n_k a valid stopping time, i.e.

$$\{n_k \le t\} \in \mathcal{H}_t, \ \forall t \in \mathbb{N}.$$

Then, any $n = n_k + s$ with some constant $s \in \mathbb{N}^+$ is also a valid stopping time because

$$\{n \le t\} = \{n_k \le t - s\} \in \mathcal{H}_{(t-s)\vee 0} \subseteq \mathcal{H}_t, \ \forall t \in \mathbb{N},$$

where $a \vee b \triangleq \max\{a, b\}$. Thus, by definition of stopping time σ-algebra from [33], we know that for any $n \ge n_k + 1$, $\mathcal{H}_n$ is defined as the collection of all sets A that have $A \cap \{n \le t\} \in \mathcal{H}_t$, $\forall t \in \mathbb{N}$[1]. Now, pick $1 \le s_1 \le s_2$ as constants, and if a set $A \in \mathcal{H}_{n_k+s_1}$, then,

$$A \cap \{n_k + s_2 \le t\} = A \cap \{n_k + s_1 \le t - (s_2 - s_1)\} \in \mathcal{H}_{(t-(s_2-s_1))\vee 0} \subseteq \mathcal{H}_t.$$

[1] An intuitive interpretation is that when $n \le t$, the set A is contained in the information known until t.

Thus, $\mathcal{H}_{n_k+s_1} \subseteq \mathcal{H}_{n_k+s_2}$ and $\{\mathcal{H}_n\}_{n\geq n_k+1}$ is indeed a filtration.

Since $\tilde{\theta}[n_k+1]$ is determined by the realization up to frame n_k, it follows, for any $t \in \mathbb{N}^+$,

$$\{\tilde{\theta}[n_k+1] \geq \theta^* + \varepsilon_0/V\} \cap \{n_k+1 \leq t\} = \cup_{s=1}^{t} \{\tilde{\theta}[s] \geq \theta^* + \varepsilon_0/V\} \in \mathcal{H}_t,$$

which implies that $\{\tilde{\theta}[n_k+1] \geq \theta^* + \varepsilon_0/V\} \in \mathcal{H}_{n_k+1}$. Since $\{\mathcal{H}_n\}_{n\geq n_k+1}$ is a filtration it follows, $\{\tilde{\theta}[n_k+1] \geq \theta^* + \varepsilon_0/V\} \in \mathcal{H}_n$ for any $n \geq n_k+1$. By the same methodology, we can show that $\{\tilde{\theta}[n] < \theta^* + \varepsilon_0/V\} \in \mathcal{H}_n$, $\forall n \geq n_k+1$, which in turn implies, $\{S_{n_k} + n_k \leq n\} \in \mathcal{H}_n$ and $\{S_{n_k} \geq n - n_k + 1\} \in \mathcal{H}_n$. Overall, the function $G[n]$ is measurable respect to $\mathcal{H}_n$, $\forall n \geq n_k+1$.

• *Proof of supermartingale claim:* It is obvious that $|G[n]| < \infty$, thus, in order to prove $G[n]$ is a supermartingale, it is enough to show that

$$\mathbb{E}\left[G[n+1] - G[n] | \mathcal{H}_n\right] \leq 0, \ \forall n \geq n_k+1. \tag{19}$$

First, on the set $\{S_{n_k} \leq n - n_k\}$, we have

$$\mathbb{E}\left[(G[n+1] - G[n])\mathbf{1}_{\{S_{n_k}+n_k\leq n\}} \middle| \mathcal{H}_n\right] = \mathbb{E}\left[(G[n] - G[n])\mathbf{1}_{\{S_{n_k}+n_k\leq n\}} \middle| \mathcal{H}_n\right] = 0.$$

It is then sufficient to show the inequality (19) holds on the set $\{S_{n_k} \geq n - n_k + 1\}$. Since

$$\begin{aligned}
&\mathbb{E}\left[G[n+1]\mathbf{1}_{\{S_{n_k}\geq n-n_k+1\}} | \mathcal{H}_n\right] \\
=&\mathbb{E}\left[\frac{e^{\lambda_{n+1}F[(n+1)\wedge(n_k+S_{n_k})]}}{\prod_{i=n_k+1}^{(n+1)\wedge(n_k+S_{n_k})}\rho_i} \middle| \mathcal{H}_n\right] \mathbf{1}_{\{\tilde{\theta}[n_k+1]\geq\theta^*+\varepsilon_0/V\}}\mathbf{1}_{\{S_{n_k}\geq n-n_k+1\}} \\
=&\mathbb{E}\left[\frac{e^{\lambda_{n+1}F[n+1]}}{\prod_{i=n_k+1}^{n+1}\rho_i} \middle| \mathcal{H}_n\right] \mathbf{1}_{\{\tilde{\theta}[n_k+1]\geq\theta^*+\varepsilon_0/V\}}\mathbf{1}_{\{S_{n_k}\geq n-n_k+1\}} \\
=&\frac{e^{\lambda_{n+1}F[n]}}{\prod_{i=n_k+1}^{n}\rho_i}\mathbb{E}\left[\frac{e^{\lambda_{n+1}(F[n+1]-F[n])}}{\rho_{n+1}} \middle| \mathcal{H}_n\right] \mathbf{1}_{\{\tilde{\theta}[n_k+1]\geq\theta^*+\varepsilon_0/V\}}\mathbf{1}_{\{S_{n_k}\geq n-n_k+1\}} \\
\leq&\frac{e^{\lambda_{n}F[n]}}{\prod_{i=n_k+1}^{n}\rho_i}\mathbb{E}\left[\frac{e^{\lambda_{n+1}(F[n+1]-F[n])}}{\rho_{n+1}} \middle| \mathcal{H}_n\right] \mathbf{1}_{\{\tilde{\theta}[n_k+1]\geq\theta^*+\varepsilon_0/V\}}\mathbf{1}_{\{S_{n_k}\geq n-n_k+1\}} \\
=&G[n]\mathbb{E}\left[\frac{e^{\lambda_{n+1}(F[n+1]-F[n])}}{\rho_{n+1}} \middle| \mathcal{H}_n\right] \mathbf{1}_{\{\tilde{\theta}[n_k+1]\geq\theta^*+\varepsilon_0/V\}}\mathbf{1}_{\{S_{n_k}\geq n-n_k+1\}},
\end{aligned}$$

where $\mathbf{1}_{\{\tilde{\theta}[n_k+1]\geq\theta^*+\varepsilon_0/V\}}$ and $\mathbf{1}_{\{S_{n_k}\geq n-n_k+1\}}$ can be moved out of the expectation because $\{\tilde{\theta}[n_k+1] \geq \theta^* + \varepsilon_0/V\} \in \mathcal{H}_n$ and $\{S_{n_k} \geq n - n_k + 1\} \in \mathcal{H}_n$, and the only inequality follows from the following argument: On the set $\{S_{n_k} \geq n - n_k + 1\}$, $\{\tilde{\theta}[n] \geq \theta^* + \varepsilon_0/V\}$, thus, by Lemma 7, $F[n] \geq 0$ and using the fact $\lambda_n > \lambda_{n+1}$, we have $\lambda_{n+1}F[n] \leq \lambda_n F[n]$. Thus, it is sufficient to show that on the set $\{S_{n_k} \geq n - n_k + 1\} \cap \{\tilde{\theta}[n_k+1] \geq \theta^* + \varepsilon_0/V\}$, we have

$$\mathbb{E}\left[\frac{e^{\lambda_{n+1}(F[n+1]-F[n])}}{\rho_{n+1}} \middle| \mathcal{H}_n\right] \leq 1.$$

By Taylor expansion, we have

$$\begin{aligned}
&\mathbb{E}\left[e^{\lambda_{n+1}(F[n+1]-F[n])} \mid \mathcal{H}_n\right] \\
=&1 + \lambda_{n+1}\mathbb{E}\left[F[n+1] - F[n] \mid \mathcal{H}_n\right] + \sum_{k=2}^{\infty}\frac{\lambda_{n+1}^k}{k!}\mathbb{E}\left[(F[n+1] - F[n])^k \mid \mathcal{H}_n\right] \\
=&1 + \lambda_{n+1}\mathbb{E}\left[F[n+1] - F[n] \mid \mathcal{H}_n\right] + \lambda_{n+1}^2\sum_{k=2}^{\infty}\frac{\lambda_{n+1}^{k-2}}{k!}\mathbb{E}\left[(F[n+1] - F[n])^k \mid \mathcal{H}_n\right] \\
\leq&1 - \frac{\lambda_{n+1}\varepsilon_0}{V} + \lambda_{n+1}^2\sum_{k=2}^{\infty}\frac{\lambda_{n+1}^{k-2}}{k!}\mathbb{E}\left[(F[n+1] - F[n])^k \mid \mathcal{H}_n\right],
\end{aligned}$$

where the last inequality comes from the following argument: On the set $\{S_{n_k} \geq n - n_k + 1\}$, $\tilde{\theta}[n_k+1] \geq \theta^* + \varepsilon_0/V$, thus, by the definition of $\tilde{\theta}[n]$, we have $\hat{\theta}[n] \geq \tilde{\theta}[n] \geq \theta^* + \varepsilon_0/V$, and Lemma 3 gives $\theta[n] \geq \theta^* + \varepsilon_0/V$, then, by Lemma 6, we have

$$\mathbb{E}\left[F[n+1] - F[n] \mid \mathcal{H}_n\right] \leq -\frac{\varepsilon_0}{V}.$$

Now, by the assumption that $V \geq \frac{\varepsilon_0 \eta}{4\log^2 2} - \frac{2\sqrt{L}}{r}$, we have $\lambda_{n+1} \leq \frac{1}{\left(\frac{2}{\eta} + \frac{4\sqrt{L}}{\eta r V}\right)\log^2(n+1)}$, which follows from simple algebraic manipulations. Using the fact that $|F[n+1] - F[n]| \leq \left(\frac{2}{\eta} + \frac{4\sqrt{L}}{\eta r V}\right)\log^2(n+1)$, we have

$$\begin{aligned}
&\mathbb{E}\left[e^{\lambda_{n+1}(F[n+1]-F[n])} \mid \mathcal{H}_n\right] \\
\leq& 1 - \frac{\lambda_{n+1}\epsilon_0}{V} + \lambda_{n+1}^2 \sum_{k=2}^{\infty} \frac{\left(\frac{1}{\left(\frac{2}{\eta} + \frac{4\sqrt{L}}{\eta r V}\right)\log^2(n+1)}\right)^{k-2}}{k!} \mathbb{E}\left[\left(\left(\frac{2}{\eta} + \frac{4\sqrt{L}}{\eta r V}\right)\log^2(n+1)\right)^k \,\middle|\, \mathcal{H}_n\right] \\
=& 1 - \frac{\lambda_{n+1}\epsilon_0}{V} + \lambda_{n+1}^2 \sum_{k=2}^{\infty} \frac{1}{k!}\left(\left(\frac{2}{\eta} + \frac{4\sqrt{L}}{\eta r V}\right)\log^2(n+1)\right)^2 \\
\leq& 1 - \frac{\lambda_{n+1}\epsilon_0}{V} + \lambda_{n+1}^2 e \left(\frac{2}{\eta} + \frac{4\sqrt{L}}{\eta r V}\right)^2 \log^4(n+1) = \rho_{n+1},
\end{aligned}$$

where the final inequality follows by completing the third term back to Taylor series which is equal to e. Overall, the inequality (19) holds and $G[n]$ is a supermartingale. □

.

Proof of Lemma 9. First of all, from Lemma 8 gives that $G[n]$ is a supermartingale starting from $n_k + 1$, thus, we have the following chains of inequalities for any $n \geq n_k + 1$:

$$\begin{aligned}
G[n_k+1] =& \mathbb{E}\left[G[n_k+1] \mid \mathcal{H}_{n_k+1}\right] \\
\geq& \mathbb{E}\left[G[n] \mid \mathcal{H}_{n_k+1}\right] \\
=& \mathbb{E}\left[\frac{e^{\lambda_n F[n\wedge(n_k+S_{n_k})]}}{\prod_{i=n_k+1}^{n}\rho_i}\mathbf{1}_{\{\tilde{\theta}[n_k+1]\geq\theta^*+\varepsilon_0/V\}} \,\middle|\, \mathcal{H}_{n_k+1}\right] \\
\geq& \mathbb{E}\left[\frac{e^{\lambda_n F[n\wedge(n_k+S_{n_k})]}}{\prod_{i=n_k+1}^{n}\rho_i}\mathbf{1}_{\{S_{n_k}\geq n-n_k+1\}}\mathbf{1}_{\{\tilde{\theta}[n_k+1]\geq\theta^*+\varepsilon_0/V\}} \,\middle|\, \mathcal{H}_{n_k+1}\right] \\
\geq& \frac{1}{\prod_{i=n_k+1}^{n}\rho_i} Pr\left[S_{n_k} \geq n - n_k + 1,\ \tilde{\theta}[n_k+1] \geq \theta^* + \varepsilon_0/V \,\middle|\, \mathcal{H}_{n_k+1}\right],
\end{aligned}$$

where the first inequality uses the supermartingale property and the last inequality uses Lemma 7 that on the set $\{S_{n_k} \geq n - n_k + 1\}$, $n \wedge (n_k + S_{n_k}) = n$ and $F[n] \geq 0$. By definition of $G[n_k+1]$,

$$G[n_k+1] = \frac{e^{\lambda_{n_k+1}F[n_k+1]}}{\rho_{n_k+1}} \leq \frac{e^{\lambda_{n_k+1}\left(\frac{2}{\eta}+\frac{4\sqrt{L}}{\eta r V}\right)\log^2(n_k+2)}}{\rho_{n_k+1}} \leq \frac{4}{3}e,$$

where the first inequality follows from the definition of $F[n]$, and the second inequality follows from the assumption that $V \geq \frac{\varepsilon_0\eta}{4\log^2 2} - \frac{2\sqrt{L}}{r}$, thus, $\lambda_{n_k+1} \leq \frac{1}{\left(\frac{2}{\eta}+\frac{4\sqrt{L}}{\eta r V}\right)\log^2(n_k+2)}$ and $\rho_{n_k+1} \geq 1 - \frac{\log^2 2}{2e} > \frac{3}{4}$. Thus, it follows,

$$Pr\left[S_{n_k} \geq n - n_k + 1,\ \tilde{\theta}[n_k+1] \geq \theta^* + \varepsilon_0/V \,\middle|\, \mathcal{H}_{n_k+1}\right] \leq \left(\prod_{i=n_k+1}^{n}\rho_i\right)\cdot\frac{4}{3}e.$$

Now, we compute

$$\begin{aligned}
&\mathbb{E}\left[S_{n_k}^4 \;\middle|\; \mathcal{H}_{n_k+1}\right] \\
=&\sum_{m=1}^{\infty} m^4 Pr\left[S_{n_k}=m \;\middle|\; \mathcal{H}_{n_k+1}\right] \\
\leq&\sum_{m=1}^{\infty}\left((m+1)^4-m^4\right) Pr\left[S_{n_k}\geq m+1,\ \tilde{\theta}[n_k+1]\geq\theta^*+\varepsilon_0/V \;\middle|\; \mathcal{H}_{n_k+1}\right]+1 \\
\leq&4\sum_{m=1}^{\infty}(m+1)^3 Pr\left[S_{n_k}\geq m+1,\ \tilde{\theta}[n_k+1]\geq\theta^*+\varepsilon_0/V \;\middle|\; \mathcal{H}_{n_k+1}\right]+1 \\
\leq&1+\frac{16}{3}e\sum_{m=1}^{\infty}(m+1)^3\prod_{i=n_k+1}^{n_k+m}\rho_i.
\end{aligned}$$

Thus, it remains to show there exists a constant C on the order $\mathcal{O}\left(V^{10}\beta^{-20}\varepsilon_0^{-10}\right)$ such that

$$\sum_{m=1}^{\infty}(m+1)^3\prod_{i=n_k+1}^{n_k+m}\rho_i\leq C(n_k+2)^{4\beta},$$

which is given is Appendix B. This implies there exists a $C_{\beta,V,\varepsilon_0}$ so that

$$\mathbb{E}\left[S_{n_k}^4\middle|\mathcal{H}_{n_k+1}\right]\leq C_{\beta,V,\varepsilon_0}(n_k+2)^{4\beta}.$$

Thus,

$$\mathbb{E}\left[S_{n_k}^4\middle|\mathcal{H}_{n_k}\right]=\mathbb{E}\left[\mathbb{E}\left[S_{n_k}^4\middle|\mathcal{H}_{n_k+1}\right]|\mathcal{H}_{n_k}\right]\leq\mathbb{E}\left[C_{\beta,V,\varepsilon_0}(n_k+2)^{4\beta}|\mathcal{H}_{n_k}\right]=C_{\beta,V,\varepsilon_0}(n_k+2)^{4\beta},$$

where the last equality follows from the fact that $n_k\in\mathcal{H}_{n_k}$. This finishes the proof. □

Appendix B: Computation of Asymptotics. In this appendix, we show that there exists a constant C such that

$$\sum_{m=1}^{\infty}(m+1)^3\prod_{i=n_k+1}^{n_k+m}\rho_i\leq Cn_k^{4\beta}.$$

We first bound ρ_i. Let $C_1=\frac{96V^2e\left(\frac{2}{\eta}+\frac{4\sqrt{L}}{\eta rV}\right)^2}{\varepsilon_0^2\beta^4}$, then,

$$\begin{aligned}
\rho_i=&1-\frac{\varepsilon_0^2}{4V^2e\left(\frac{2}{\eta}+\frac{4\sqrt{L}}{\eta rV}\right)^2\log^4(i+1)} \\
=&1-\frac{1}{C_1\frac{\beta^4}{24}\log^4(i+1)} \\
<&1-\frac{1}{C_1(i+1)^{\beta}},
\end{aligned}$$

where we used the fact that $\frac{\beta^4}{24}\log^4(i+1)<(i+1)^{\beta},\ \forall\beta>0,i\geq0$. Next, to bound $\prod_{i=n_k+1}^{n_k+m}\rho_i$, we take the logarithm:

$$\log\left(\prod_{i=n_k+1}^{n_k+m}\rho_i\right)=\sum_{i=n_k+1}^{n_k+m}\log\rho_i$$

$$
\begin{aligned}
&= \sum_{i=n_k+1}^{n_k+m} \log\left(1 - \frac{1}{C_1(i+1)^\beta}\right) \\
&\le -\sum_{i=n_k+1}^{n_k+m} \frac{1}{C_1(i+1)^\beta} \\
&\le -\frac{1}{C_1}\int_{n_k+2}^{n_k+m+1} \frac{1}{x^\beta} dx.
\end{aligned}
$$

where the first inequality follows from the first order Taylor expansion. Since $\beta < 1$, we compute the integral, which gives

$$
-\frac{1}{C_1}\int_{n_k+2}^{n_k+m+1} \frac{1}{x^\beta} dx = -\frac{1}{C_1(1-2\beta)}\left((n_k+m+1)^{1-\beta} - (n_k+2)^{1-\beta}\right).
$$

Thus,

$$
\begin{aligned}
&\sum_{m=1}^{\infty}(m+1)^3 \prod_{i=n_k+1}^{n_k+m} \rho_i \\
&\le \sum_{m=1}^{\infty}(m+1)^3 e^{-\frac{1}{C_1(1-\beta)}\left((n_k+m+1)^{1-\beta} - (n_k+2)^{1-\beta}\right)} \\
&\le \int_0^{\infty}(x+2)^3 e^{-\frac{1}{C_1(1-\beta)}\left((x+n_k+2)^{1-\beta} - (n_k+2)^{1-\beta}\right)} dx + (3C_1(1-\beta))^4,
\end{aligned}
$$

where the last inequality follows from the fact that the integrand is monotonically decreasing when $x > 3C_1(1-\beta)$, thus, the integral dominates the sum on the tail $x > 3C_1(1-\beta)$. For the part where $x \le 3C_1(1-\beta)$, the maximum of the integrand is bounded by $(3C_1(1-\beta))^3$. Thus, the total difference of such approximation is bounded by $(3C_1(1-\beta))^4$. Then, we try to estimate the integral. Notice that

$$
\frac{d}{dx} e^{-\frac{1}{C_1(1-\beta)}(x+n_k+2)^{1-\beta}} = -\frac{1}{C_1} e^{-\frac{1}{C_1(1-\beta)}(x+n_k+2)^{1-\beta}} (x+n_k+2)^{-\beta},
$$

we do integration-by-parts, which gives

$$
\begin{aligned}
&\int_0^{\infty}(x+2)^3 e^{-\frac{1}{C_1(1-\beta)}\left((x+n_k+2)^{1-\beta} - (n_k+2)^{1-\beta}\right)} dx \\
=&\int_0^{\infty}(x+2)^3 (x+n_k+2)^{\beta}(x+n_k+2)^{-\beta} e^{-\frac{1}{C_1(1-\beta)}(x+n_k+2)^{1-\beta}} dx \cdot e^{\frac{1}{C_1(1-\beta)}(n_k+2)^{1-\beta}} \\
=&8C_1(n_k+2)^{\beta} + \int_0^{\infty} C_1\left(3(x+2)^2(x+n_k+2)^{\beta} + \beta(x+2)^3(x+n_k+2)^{\beta-1}\right) e^{-\frac{1}{C_1(1-\beta)}\left((x+n_k+2)^{1-\beta} - (n_k+2)^{1-\beta}\right)} dx.
\end{aligned}
$$

Since $5\beta \le 1$ and $n_k \ge 1$, we have $x + n_k + 2 \ge x + 2$, which implies $(x+2)^3(x+n_k+2)^{\beta-1} \le (x+2)^2(x+n_k+2)^{\beta}$, thus,

$$
\begin{aligned}
&\int_0^{\infty}(x+2)^3 e^{-\frac{1}{C_1(1-\beta)}\left((x+n_k+2)^{1-\beta} - (n_k+2)^{1-\beta}\right)} dx \\
\le&8C_1(n_k+2)^{\beta} + \int_0^{\infty} 4C_1(x+2)^2(x+n_k+2)^{\beta} e^{-\frac{1}{C_1(1-\beta)}\left((x+n_k+2)^{1-\beta} - (n_k+2)^{1-\beta}\right)} dx.
\end{aligned}
$$

Repeat above procedure 3 more times, we have

$$
\begin{aligned}
&\int_0^{\infty}(x+2)^3 e^{-\frac{1}{C_1(1-\beta)}\left((x+n_k+2)^{1-\beta} - (n_k+2)^{1-\beta}\right)} dx \\
\le&8C_1(n_k+2)^{\beta} + 16C_1^2(n_k+2)^{2\beta} + 24C_1^3(n_k+2)^{3\beta} + 24C_1^4(n_k+2)^{4\beta} \\
&+ \int_0^{\infty} 24C_1^4(x+n_k+2)^{4\beta-1} e^{-\frac{1}{C_1(1-\beta)}\left((x+n_k+2)^{1-\beta} - (n_k+2)^{1-\beta}\right)} dx \\
\le&8C_1(n_k+2)^{\beta} + 16C_1^2(n_k+2)^{2\beta} + 24C_1^3(n_k+2)^{3\beta} + 24C_1^4(n_k+2)^{4\beta} + 24C_1^5 \le C(n_k+2)^{5\beta},
\end{aligned}
$$

for some C on the order of C_1^5 (which is $\mathcal{O}\left(V^{10}\beta^{-20}\varepsilon_0^{-10}\right)$), where the second to last inequality follows from $4\beta-1\leq-\beta$ and thus, we replace $(x+n_k+2)^{4\beta-1}$ with $(x+n_k+2)^{-\beta}$ and do a direct integration. Overall, we proved the claim.